Networking Concepts for the Accelar 1000 Series Routing Switch

Software Release 2.0

Part No. 205588-A
March 1999

NORTEL NETWORKS™

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Copyright © 1999 Bay Networks, Inc.

All rights reserved. Printed in the USA. March 1999.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Bay Networks, Inc.

The software described in this document is furnished under a license agreement and may only be used in accordance with the terms of that license. A summary of the Software License is included in this document.

Trademarks

Bay Networks is a registered trademark of Bay Networks, Inc.

Accelar, BayStack, LinkSafe, and the Bay Networks logo are trademarks of Bay Networks, Inc.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective owners.

Restricted Rights Legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Bay Networks, Inc. reserves the right to make changes to the products described in this document without notice.

Bay Networks, Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product are Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Bay Networks, Inc. Software License Agreement

NOTICE: Please carefully read this license agreement before copying or using the accompanying software or installing the hardware unit with pre-enabled software (each of which is referred to as "Software" in this Agreement). BY COPYING OR USING THE SOFTWARE, YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. THE TERMS EXPRESSED IN THIS AGREEMENT ARE THE ONLY TERMS UNDER WHICH BAY NETWORKS WILL PERMIT YOU TO USE THE SOFTWARE. If you do not accept these terms and conditions, return the product, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

1. License Grant. Bay Networks, Inc. ("Bay Networks") grants the end user of the Software ("Licensee") a personal, nonexclusive, nontransferable license: (a) to use the Software either on a single computer or, if applicable, on a single authorized device identified by host ID, for which it was originally acquired; (b) to copy the Software solely for backup purposes in support of authorized use of the Software; and (c) to use and copy the associated user manual solely in support of authorized use of the Software by Licensee. This license applies to the Software only and does not extend to Bay Networks Agent software or other Bay Networks software products. Bay Networks Agent software or other Bay Networks software products are licensed for use under the terms of the applicable Bay Networks, Inc. Software License Agreement that accompanies such software and upon payment by the end user of the applicable license fees for such software.

2. Restrictions on use; reservation of rights. The Software and user manuals are protected under copyright laws. Bay Networks and/or its licensors retain all title and ownership in both the Software and user manuals, including any revisions made by Bay Networks or its licensors. The copyright notice must be reproduced and included with any copy of any portion of the Software or user manuals. Licensee may not modify, translate, decompile, disassemble, use for any competitive analysis, reverse engineer, distribute, or create derivative works from the Software or user manuals or any copy, in whole or in part. Except as expressly provided in this Agreement, Licensee may not copy or transfer the Software or user manuals, in whole or in part. The Software and user manuals embody Bay Networks' and its licensors' confidential and proprietary intellectual property. Licensee shall not sublicense, assign, or otherwise disclose to any third party the Software, or any information about the operation, design, performance, or implementation of the Software and user manuals that is confidential to Bay Networks and its licensors; however, Licensee may grant permission to its consultants, subcontractors, and agents to use the Software at Licensee's facility, provided they have agreed to use the Software only in accordance with the terms of this license.

3. Limited warranty. Bay Networks warrants each item of Software, as delivered by Bay Networks and properly installed and operated on Bay Networks hardware or other equipment it is originally licensed for, to function substantially as described in its accompanying user manual during its warranty period, which begins on the date Software is first shipped to Licensee. If any item of Software fails to so function during its warranty period, as the sole remedy Bay Networks will at its discretion provide a suitable fix, patch, or workaround for the problem that may be included in a future Software release. Bay Networks further warrants to Licensee that the media on which the Software is provided will be free from defects in materials and workmanship under normal use for a period of 90 days from the date Software is first shipped to Licensee. Bay Networks will replace defective media at no charge if it is returned to Bay Networks during the warranty period along with proof of the date of shipment. This warranty does not apply if the media has been damaged as a result of accident, misuse, or abuse. The Licensee assumes all responsibility for selection of the Software to achieve Licensee's intended results and for the installation, use, and results obtained from the Software. Bay Networks does not warrant (a) that the functions contained in the software will meet the Licensee's requirements, (b) that the Software will operate in the hardware or software combinations that the Licensee may select, (c) that the operation of the Software will be uninterrupted or error free, or (d) that all defects in the operation of the Software will be corrected. Bay Networks is not obligated to remedy any Software defect that cannot be reproduced with the latest Software release. These warranties do not apply to the Software if it has been (i) altered, except by Bay Networks or in accordance with its instructions; (ii) used in conjunction with another vendor's product, resulting in the defect; or (iii) damaged by improper environment, abuse, misuse, accident, or negligence. THE FOREGOING WARRANTIES AND LIMITATIONS ARE EXCLUSIVE REMEDIES AND ARE IN LIEU OF ALL OTHER WARRANTIES EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Licensee is responsible for the security of its own data and information and for maintaining adequate procedures apart from the Software to reconstruct lost or altered files, data, or programs.

4. Limitation of liability. IN NO EVENT WILL BAY NETWORKS OR ITS LICENSORS BE LIABLE FOR ANY COST OF SUBSTITUTE PROCUREMENT; SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES; OR ANY DAMAGES RESULTING FROM INACCURATE OR LOST DATA OR LOSS OF USE OR PROFITS ARISING OUT OF OR IN CONNECTION WITH THE PERFORMANCE OF THE SOFTWARE, EVEN IF BAY NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF BAY NETWORKS RELATING TO THE SOFTWARE OR THIS AGREEMENT EXCEED THE PRICE PAID TO BAY NETWORKS FOR THE SOFTWARE LICENSE.

5. Government Licensees. This provision applies to all Software and documentation acquired directly or indirectly by or on behalf of the United States Government. The Software and documentation are commercial products, licensed on the open market at market prices, and were developed entirely at private expense and without the use of any U.S. Government funds. The license to the U.S. Government is granted only with restricted rights, and use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1) of the Commercial Computer Software—Restricted Rights clause of FAR 52.227-19 and the limitations set out in this license for civilian agencies, and subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013, for agencies of the Department of Defense or their successors, whichever is applicable.

6. Use of Software in the European Community. This provision applies to all Software acquired for use within the European Community. If Licensee uses the Software within a country in the European Community, the Software Directive enacted by the Council of European Communities Directive dated 14 May, 1991, will apply to the examination of the Software to facilitate interoperability. Licensee agrees to notify Bay Networks of any such intended examination of the Software and may procure support and assistance from Bay Networks.

7. Term and termination. This license is effective until terminated; however, all of the restrictions with respect to Bay Networks' copyright in the Software and user manuals will cease being effective at the date of expiration of the Bay Networks copyright; those restrictions relating to use and disclosure of Bay Networks' confidential information shall continue in effect. Licensee may terminate this license at any time. The license will automatically terminate if Licensee fails to comply with any of the terms and conditions of the license. Upon termination for any reason, Licensee will immediately destroy or return to Bay Networks the Software, user manuals, and all copies. Bay Networks is not liable to Licensee for damages in any form solely by reason of the termination of this license.

8. Export and Re-export. Licensee agrees not to export, directly or indirectly, the Software or related technical data or information without first obtaining any required export licenses or other governmental approvals. Without limiting the foregoing, Licensee, on behalf of itself and its subsidiaries and affiliates, agrees that it will not, without first obtaining all export licenses and approvals required by the U.S. Government: (i) export, re-export, transfer, or divert any such Software or technical data, or any direct product thereof, to any country to which such exports or re-exports are restricted or embargoed under United States export control laws and regulations, or to any national or resident of such restricted or embargoed countries; or (ii) provide the Software or related technical data or information to any military end user or for any military end use, including the design, development, or production of any chemical, nuclear, or biological weapons.

9. General. If any provision of this Agreement is held to be invalid or unenforceable by a court of competent jurisdiction, the remainder of the provisions of this Agreement shall remain in full force and effect. This Agreement will be governed by the laws of the state of California.

Should you have any questions concerning this Agreement, contact Bay Networks, Inc., 4401 Great America Parkway, P.O. Box 58185, Santa Clara, California 95052-8185.

LICENSEE ACKNOWLEDGES THAT LICENSEE HAS READ THIS AGREEMENT, UNDERSTANDS IT, AND AGREES TO BE BOUND BY ITS TERMS AND CONDITIONS. LICENSEE FURTHER AGREES THAT THIS AGREEMENT IS THE ENTIRE AND EXCLUSIVE AGREEMENT BETWEEN BAY NETWORKS AND LICENSEE, WHICH SUPERSEDES ALL PRIOR ORAL AND WRITTEN AGREEMENTS AND COMMUNICATIONS BETWEEN THE PARTIES PERTAINING TO THE SUBJECT MATTER OF THIS AGREEMENT. NO DIFFERENT OR ADDITIONAL TERMS WILL BE ENFORCEABLE AGAINST BAY NETWORKS UNLESS BAY NETWORKS GIVES ITS EXPRESS WRITTEN CONSENT, INCLUDING AN EXPRESS WAIVER OF THE TERMS OF THIS AGREEMENT.

Contents

Preface

Before You Begin xv
Related Publications xvii
How to Get Help xx

Chapter 1
Introduction

Accelar Hardware and Software 1-1
Compatibility 1-2
Concepts 1-4

Chapter 2
Accelar Management Basics

Boot Sequence 2-1
Stage 1: Boot Monitor Image Load 2-2
Stage 2: Boot Configuration Load 2-2
Stage 3: Run-Time Image Load 2-2
Stage 4: Routing Switch Configuration Load 2-3
Flash/PCMCIA File System 2-4
Flash Memory Organization 2-4
Boot Flash 2-4
System Flash (flash:) 2-5
PCMCIA (pcmcia:) 2-6
File Types 2-6
Executables 2-6
Log Files 2-7
Configuration Files 2-7
Script Files 2-7
Trace Logs 2-7

Management Tools 2-8
Accelar Device Manager 2-8
Accelar VLAN Manager 2-8
Boot Monitor Command Line Interface (CLI) 2-9
Run-Time Command Line Interface (CLI) 2-9
Accelar Configuration Page 2-10
Accelar Access Levels and Passwords 2-11
Read-Only Access 2-11
Layer 2 Read-Write Access 2-11
Layer 3 Read-Write Access 2-11
Read-Write Access 2-11
Read-Write-All Access 2-11
CLI Access 2-12
SNMP Community Strings 2-12
Web Parameters 2-13
Management Access 2-13
Port Lock 2-14

Chapter 3
Layer 2 Networking Concepts

VLANs 3-1
Port-Based VLANs 3-2
Policy-Based VLANs 3-3
Source MAC-Based VLANs 3-4
Protocol-Based VLANs 3-5
Source IP Subnet-Based VLANs 3-7
VLAN Tagging and Port Types 3-8
802.1Q Tagged Ports 3-9
Explicit Tagging Versus Implicit Tagging 3-9
IP Routing and VLANs 3-10
Spanning Tree Protocol, Groups, and FastStart 3-10
Spanning Tree Protocol 3-11
Accelar Spanning Tree Groups 3-11
Accelar Spanning Tree FastStart 3-12

Traffic Prioritization 3-12
Setting Priority 3-13
Accelar 1000 Series VLANs 3-14
VLAN Rules 3-14
Special VLANs 3-15
Default VLAN 3-16
Unassigned VLAN 3-16
Brouter Ports 3-16
Default Configuration 3-17
Multi-Link Trunking 3-17
Multi-Link Trunking Examples 3-18
Client/Server Configuration Utilizing Multi-Link Trunks 3-20
Network Management and Diagnostics 3-22
RMON 3-22
Port Mirroring 3-23
Syslog 3-24

Chapter 4
IP Interfaces and Router Management

IP Addresses 4-1
Subnet Addressing 4-2
Supernet Addressing 4-2
Types of IP Routing 4-3
Isolated Routing Ports 4-3
Virtual Routing Between VLANs 4-4
Brouter Ports 4-5
Static Routes 4-6
Router Management 4-6
Address Resolution Protocol (ARP) 4-6
Using Proxy ARP 4-7
Flushing Router Tables 4-8
BootP/DHCP Relay 4-8
Differences Between DHCP and BootP 4-8
Summary of DHCP Relay Operation 4-8
Forwarding DHCP Packets 4-9
Multiple BootP/DHCP Servers 4-10

UDP Broadcast Forwarding 4-11
UDP Forwarding Operation 4-11
Reverse Address Resolution Protocol (RARP) 4-12
Virtual Router Redundancy Protocol (VRRP) 4-12
Dynamic IP Routing Protocols 4-15
Routing Information Protocol (RIP) 4-15
Open Shortest Path First (OSPF) Protocol 4-16
OSPF Addresses and Variable-Length Masks 4-17
OSPF Neighbors 4-17
Neighbor Adjacencies 4-18
Designated Routers 4-18
OSPF Areas 4-18
OSPF Router Types 4-19
AS External Routes 4-20
IP Policies 4-20
Policies in Accelar Switches 4-21

Chapter 5
IP Multicasting

Multicast Host Groups 5-1
Multicast Addresses 5-2
Internet Group Management Protocol 5-3
IGMP Concepts and Terminology 5-3
IGMP Queries 5-3
IGMP Host Reports 5-4
Host Leave Messages 5-4
Accelar IGMP Implementation 5-5
IGMP Snooping 5-5
IGMP Proxy 5-6
Distance Vector Multicast Routing Protocol 5-6

DVMRP Concepts and Terminology 5-7
Neighbor Connections 5-8
Source Route Advertisements 5-8
How DVMRP Chooses a Route 5-9
Routing Table 5-9
Shortest-Path Trees 5-10
Accelar Implementation of DVRMP 5-10

Chapter 6
IP Filtering

Filter Characteristics 6-1
Source and Destination Filters 6-2
Global Filters 6-3
Filter Configuration 6-3
Actions 6-4

Chapter 7
IPX Routing

IPX Protocol 7-1
IPX Packets 7-4
IPX Traffic 7-6
Accelar IPX Network-Layer Support 7-7
Routing Information Protocol (RIP) 7-7
Service Advertising Protocol (SAP) 7-9
Broadcast Techniques 7-10
Best Route Algorithm 7-10
Split Horizon 7-11
NetBIOS 7-12
Static Routes 7-14
Static SAP Services 7-14
IPX Default Static Routes 7-14

Appendix A
Port Numbering and MAC Address Assignment

Port Numbering A-1
MAC Address Assignment A-3
Base MAC Address A-3
Physical MAC Addresses A-4
Virtual MAC Addresses A-5

Index

Figures

Figure 3-1. Example of Port-Based VLANs 3-3
Figure 3-2. Example of a Dynamic VLAN Based on Protocol 3-5
Figure 3-3. Example of IEEE 802.1Q Tagged Frame Format 3-8
Figure 3-4. Example of Explicit Encapsulation Tagging 3-8
Figure 3-5. Multiple Spanning Tree Groups 3-11
Figure 3-6. Switch-to-Switch MLT Configuration Example 3-19
Figure 3-7. Switch-to-Server MLT Configuration Example 3-20
Figure 3-8. Client/Server Configuration Example 3-21

Figure 4-1. Routing Between IP Destination Addresses 4-4
Figure 4-2. IP Routing Between VLANs 4-5
Figure 4-3. Proxy ARP Operation 4-7
Figure 4-4. Example of DHCP Operation 4-9
Figure 4-5. Forwarding DHCP Packets 4-9
Figure 4-6. Configuring Multiple BootP/DHCP Servers 4-10
Figure 4-7. Example of VRRP in a Network 4-13
Figure 4-8. Hop Count or Metric in RIP 4-16

Figure 7-1. IPX Internetwork 7-3
Figure 7-2. IPX Header 7-5
Figure 7-3. IPX Packets 7-5
Figure 7-4. IPX Addressing 7-6
Figure 7-5. IPX-RIP Packet 7-8
Figure 7-6. IPX-SAP Packet 7-10
Figure 7-7. Best Route Algorithm Example 7-11
Figure 7-8. Split Horizon Enabled 7-12
Figure 7-9. NetBIOS Support in a NetWare Environment 7-13

Figure A-1. Accelar 1200 Slots A-1
Figure A-2. Accelar 1l00 Slots A-2
Figure A-3. Port Numbering on I/O Modules A-2
Figure A-4. Slot and Port Numbering on the Accelar 1050/1051 Switch A-3

Tables

Table 3-1. Reserved PIDs for User-Defined Protocol-Based VLANs 3-6

Table 4-1. Router Classifications 4-19

Table 5-1. Parts of a Routing Table Entry 5-9

Table 6-1. Port Actions for Combinations of Matching Filters 6-4

Table A-1. Last Byte of Physical MAC Address A-5

Preface

Welcome to the Bay Networks® Accelar™ 1000 Series Routing Switches, a family of switches that enable integration of switching (layer 2) and routing (layer 3) functions in a single device.

This guide summarizes the general networking concepts used in the Accelar 1000 Series routing switches. These features can be implemented using the Accelar Management Software Graphical User Interface (GUI) or the command line interface (CLI). For detailed information about implementing these functions in the Device Manager GUI, refer to *Reference for Accelar Management Software Switching Operations* and *Reference for Accelar Management Software Routing Operations*. For implementation using the CLI, refer to *Reference for the Accelar 1000 Series Command Line Interface*.

Before You Begin

This guide is intended for network administrators with the following background:

- Basic knowledge of networks, Ethernet bridging, IP routing, and IPX routing
- Familiarity with networking concepts and terminology
- Basic knowledge of network topologies

Text Conventions

This guide uses the following text conventions:

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: If the command syntax is: **ping** *<ip_address>*, you enter: **ping 192.32.10.12**
bold text	Indicates text that you need to enter and command names and options. Example: Enter **show ip** {**alerts** | **routes**}
braces ({})	Indicate required elements in syntax descriptions where there is more than one option. You must choose only one of the options. Do not type the braces when entering the command. Example: If the command syntax is: **show ip** {**alerts** | **routes**}, you must enter either: **show ip alerts** or **show ip routes**.
brackets ([ ])	Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command. Example: If the command syntax is: **show ip interfaces** [**-alerts**], you can enter either: **show ip interfaces** or **show ip interfaces -alerts**.
italic text	Indicates file and directory names, new terms, and book titles.
`screen text`	Indicates system output, for example, prompts and system messages. Example: `Set Bay Networks Trap Monitor Filters`

separator (>)	Shows menu paths. Example: Protocols > IP identifies the IP option on the Protocols menu.
vertical line (|)	Separates choices for command keywords and arguments. Enter only one of the choices. Do not type the vertical line when entering the command. Example: If the command syntax is: **show ip** {**alerts** | **routes**}, you enter either: **show ip alerts** or **show ip routes**, but not both.

Related Publications

For more information about using Accelar Management Software or Accelar routing switches, refer to the following publications:

- *Installing the Accelar 1000 Series Chassis*
  (Bay Networks part number 893-01051-D)

  Outlines the procedures for installing and booting your Accelar routing switch, as well as instructions for installing the Accelar Management Software.

- *Using the Accelar 1200/1250 Routing Switch*
  (Bay Networks part number 893-01049-C)

  Provides information about Accelar 1200 and Accelar 1250 switches, including operating specifications and common procedures.

- *Using the Accelar 1100/1150 Routing Switch*
  (Bay Networks part number 893-01050-C)

  Provides information about Accelar 1l00, Accelar 1100R, Accelar 1150, and Accelar 1150R switches, including operating specifications and common procedures.

- *Using the Accelar 1050/1051 Routing Switch*
  (Bay Networks part number 201603-C)

  Provides information about Accelar 1050 and 1051 routing switches, including operating specifications and common procedures.

- *Reference for Accelar Management Software Switching Operations* (Bay Networks part number 205586-A)

  Describes how to use Device Manager to configure and manage layer 2 (switching) functions with the Accelar routing switch, including procedures and illustrations of pertinent screens.

- *Reference for Accelar Management Software Routing Operations* (Bay Networks part number 205587-A)

  Describes how to use Device Manager to configure and manage layer 3 (routing) functions with the Accelar routing switch, including procedures and illustrations of pertinent screens.

- *Reference for the Accelar 1000 Series Command Line Interface* (Bay Networks part number 202086-B)

  Describes the command line interface (CLI) commands and parameters. Most configuration tasks that can be performed using Device Manager and VLAN Manager can also be done using the CLI.

- *Release Notes for the Accelar 1000 Series Products Software Release 2.0* (Bay Networks part number 896-00181-E)

  Documents important changes about the software or hardware that are not covered in other related publications.

For more information about networking concepts, protocols, and topologies, you may want to consult the following sources:

- RFC 1058 (RIP version 1)
- RFC 1723 (RIP version 2)
- RFC 1213 (IP)
- RFC 1389 (RIP 2 Management Information Base)
- RFC 1493 (Bridge MIB)
- RFC 1573 (IANAIf Type)
- RFC 1643 (Ether-like MIB)
- RFC 1757 (RMON)
- RFC 1271 (RMON)

- RFC 1850 (OSPF MIB)
- RFC 1253 (OSPF)
- RFC 1583 (OSPF)
- RFC 2178 (OSPF)
- IEEE 802.1D (Standard for Spanning Tree Protocol)
- IEEE 802.3 (Ethernet)
- IEEE 802.1Q (VLAN Tagging)
- Enterprise MIB (located on the *Accelar 1000 Series Software CD*)

You can now print Bay Networks technical manuals and release notes free, directly from the Internet. Go to *support.baynetworks.com/library/tpubs/*. Find the Bay Networks product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Using Adobe Acrobat Reader, you can open the manuals and release notes, search for the sections you need, and print them on most standard printers. You can download Acrobat Reader free from the Adobe Systems Web site, *www.adobe.com*.

You can purchase Bay Networks documentation sets, CDs, and selected technical publications through the Bay Networks Collateral Catalog. The catalog is located on the World Wide Web at *support.baynetworks.com/catalog.html* and is divided into sections arranged alphabetically:

- The “CD ROMs” section lists available CDs.
- The “Guides/Books” section lists books on technical topics.
- The “Technical Manuals” section lists available printed documentation sets.

Make a note of the part numbers and prices of the items that you want to order. Use the “Marketing Collateral Catalog description” link to place an order and to print the order form.

How to Get Help

If you purchased a service contract for your Bay Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Bay Networks service program, contact one of the following Bay Networks Technical Solutions Centers:

Technical Solutions Center	Telephone Number	Fax Number
Billerica, MA	800-2LANWAN (800-252-6926)	978-916-3514
Santa Clara, CA	800-2LANWAN (800-252-6926)	408-495-1188
Valbonne, France	33-4-92-96-69-68	33-4-92-96-69-98
Sydney, Australia	61-2-9927-8800	61-2-9927-8811
Tokyo, Japan	81-3-5402-0180	81-3-5402-0173

Chapter 1
Introduction

The Accelar 1000 Series routing switches offer great flexibility of use in a network design because they combine bridging/switching, VLANs, IP routing, and IPX routing capabilities in a single device. The integration of switching (layer 2) and routing (layer 3) functions allows users to flexibly deploy routing and switching technology as needed and to build high-performance, fully switched networks while implementing traditional routing services. Accelar switches are capable of low latency routing—routing functionality at switching speeds. Using the Accelar routing switch to build large switched networks allows you to offload IP routing from current backbone routers and to provide much higher throughput.

This guide summarizes the general networking concepts used in the Accelar 1000 Series routing switches. For detailed information about implementing these functions using the Device Manager GUI, refer to *Reference for Accelar Management Software Switching Operations* and *Reference for Accelar Management Software Routing Operations*. For implementation using the command line interface (CLI), refer to *Reference for the Accelar 1000 Series Command Line Interface*.

Accelar Hardware and Software

Accelar switches support 10/100 megabits per second (Mb/s) Ethernet/Fast Ethernet as well as Gigabit Ethernet technology and include the following:

- Accelar 1200 eight-slot chassis
- Accelar 1250 four-slot chassis
- Accelar 1100 16-port 10/100BASE-TX standalone configuration with two expansion module slots

- Accelar 1150 4-port 1000BASE-SX Gigabit Ethernet standalone configuration with two expansion module slots
- Accelar 1050 or 1051 fixed configuration with 12 ports of 10/100BASE-TX Ethernet and one port of 1000BASE-SX Gigabit Ethernet, with the Accelar 1051 switch including LinkSafe™ redundant Gigabit links

Expansion modules for the Accelar 1200/1250 and 1100/1150 switches provide a variety of speed and port configuration options, including 1000BASE-SX, 1000BASE-LX, and very long distance Gigabit Ethernet modules.

Accelar routing switches provide aggregate forwarding capacity up to 7 million packets per second (Mpps). Supported protocols include IP, IPX, RIP, RIP2, OSPF, IGMP, and DVMRP. The Accelar 1200 switch offers the option of installing a second Silicon Switch Fabric (SSF) module for CPU redundancy. Accelar 1200, 1100, and 1150 switches accept redundant power supplies. For protection against cable faults, the Accelar 1200, 1250, 1100, and 1051 switches allow installation of LinkSafe redundant link modules on Gigabit Ethernet ports.

Software management is provided through the Device Manager and VLAN Manager graphical user interfaces (GUIs), as well as a complete command line Interface (CLI). The software operates on Sun Microsystems Solaris, HP-UX, and IBM AIX workstations, as well as PCs running Windows 95, Windows 98, or Windows NT.

Compatibility

Some features in each software release require certain levels of supporting hardware in order to function. The version -A hardware includes ARU2 ASICs, and the version -B hardware includes ARU3 ASICs. Major software release features are listed below with hardware dependencies noted.

Release 1.3:

- CLI enhancements
- Virtual Router Redundancy Protocol (VRRP)
- Accept and announce route policies
- IP Multicast optimization at layer 2
- Access and security features
- RIP, OSPF, and VLAN enhancements

- Multi-Link Trunking (requires -A version hardware)
- IP prefix flow filters (requires -A version hardware)
- IGMPv1 snooping (requires -A version hardware)
- Syslog

Release 2.0:

- Internetwork Packet Exchange (IPX) Protocol (requires -B version hardware)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Internet Group Management Protocol (IGMP) version 2 (requires -B version hardware)
- VRRP, IGMP, and IP flow filter enhancements
- UDP forwarding (NetBIOS)
- Reverse Address Resolution Protocol (RARP)
- Brouter interfaces
- Discard unknown MAC security
- OSPF, VLAN, and CLI enhancements

Some advanced features in existing functionality may also require -B version hardware.

To determine the hardware version(s) in your chassis, you can verify the ASIC versions in your module or chassis.

- In the CLI, enter the command: `show system info.` The resulting display will indicate the ARU level of the chassis and, if applicable, the cards.
- In Device Manager, choose Edit > Chassis > Chassis. The ARU and QUID modes are listed in the display screen.

Concepts

Devising a network design that optimizes the features in an Accelar routing switch requires an understanding of some advanced networking concepts.This guide discusses advanced networking concepts with a special emphasis on how these concepts are implemented in an Accelar routing switch. This manual assumes that you are familiar with the basics of bridging/switching, routing, and Spanning Tree Protocol.

The following information is included in this guide:

- Chapter 2 includes general information about Accelar switching software including boot and flash memory, description of tools, and security information.
- Chapter 3 includes layer 2 (switching) concepts as used in Accelar software, including how the switch handles VLANs, Spanning Tree Protocol, Multi-Link Trunking, and network management.
- Chapter 4 describes layer 3 (routing) concepts, including IP addressing, types of IP routing, router versus brouter ports, ARP, DHCP, UDP broadcast forwarding, RARP, VRRP, RIP, and OSPF. It also describes how IP announce and accept policies can be used for RIP and OSPF to control the flow of data to and from the routing tables.
- Chapter 5 describes the IP multicasting support in Accelar switches—IGMP and DVMRP.
- Chapter 6 describes how Accelar switches implement IP filtering for traffic management.
- Chapter 7 describes the Accelar software support for IPX routing.

Chapter 2
Accelar Management Basics

This chapter describes the basic operational process in the Accelar 1000 Series routing switch. Topics covered in this chapter include the following information:

- Boot sequence (this page)
- Flash/PCMCIA file system (page 2-4)
- Overview of management tools (page 2-8)
- Accelar access levels and passwords (page 2-11)

Boot Sequence

Accelar 1000 Series routing switches go through a four-stage boot sequence before becoming fully functional routing switches.

The boot sequence includes the following four stages:

1. Boot monitor image load (page 2-2)
2. Boot configuration load (page 2-2)
3. Run-time image load (page 2-2)
4. Routing switch configuration load (page 2-3)

The following sections describe what happens at each stage in the boot process.

Stage 1: Boot Monitor Image Load

At the power-up or reset sequence, the Silicon Switch Fabric (SSF) module loads the boot monitor image. The boot monitor image is contained in flash memory on the SSF module. If an Accelar 1200 routing switch contains a redundant SSF module, the first SSF module installed or the module in slot 4 becomes the active SSF module on powering up or resetting. Consequently, the boot monitor image is loaded from the flash memory on that SSF module.

When the boot monitor image is loaded, the CPU and basic system devices such as the console port, modem port, PCMCIA card slot (if applicable), and debug Ethernet port are initialized. Note that the I/O ports are not available at this stage. The I/O ports are not initialized until later in the boot process.

Stage 2: Boot Configuration Load

After the bootstrap image loads, the boot configuration is loaded. The boot configuration resides in boot flash memory on the SSF module, and it consists of parameters that control how the boot process proceeds and how the devices initialized by the boot monitor are configured. For information about boot monitor commands, refer to the section on the Boot Monitor CLI in *Reference for the Accelar 1000 Series Command Line Interface*.

If Autoboot is disabled or interrupted at the console, the boot process stops. At this stage, the user has access to the Boot Monitor CLI at the console.

Stage 3: Run-Time Image Load

The run-time image loads after the boot configuration. This software image initializes the I/O modules and provides full routing switch functionality.

The run-time image can be loaded from various sources depending on the Accelar switch model:

- Accelar 1200/1250 switches can load the run-time image from the flash memory, from a PCMCIA card, or from a TFTP server using the diagnostic Ethernet port.
- Accelar 1100/1150 switches can load the run-time image from the flash memory or from a TFTP server using the diagnostic Ethernet port.
- Accelar 1050/1051 switches load the run-time image from the flash memory.

The factory default load order is: PCMCIA (if applicable), flash memory, and TFTP. However, you can define the source and order from which to load the run-time image:

- To specify the order in the Boot Monitor CLI, refer to the choices command in the Boot Monitor CLI section of *Reference for the Accelar 1000 Series Command Line Interface*.
- To specify the order in Accelar Device Manager, go to the Edit > Chassis > Boot screen.
- To specify the source using the Run-Time CLI commands, refer to the section on File and Device Management of *Reference for the Accelar 1000 Series Command Line Interface*.

Stage 4: Routing Switch Configuration Load

The final step before the boot process is complete is to load the routing switch configuration. The routing switch configuration consists of any higher-level functionality, including:

- Chassis configuration
- Port configuration
- Spanning tree group configuration
- VLAN configuration
- Routing configuration
- IP address assignments
- RMON configuration

The default configuration includes:

- A single, port-based default VLAN with a VLAN identification number of 1, bound to the default spanning tree group.
- All ports in a single spanning tree group, STG number 1. The default spanning tree group is 802.1D compliant, and its BPDUs are never tagged.
- Spanning Tree FastStart disabled on all ports.
- No interfaces assigned IP addresses.
- Traffic priority for all ports set to normal priority.

- All ports as nontagged ports. Untagged frames are tagged on egress out a tagged port, and tagged frames are untagged on egress out a nontagged port.

Whether the routing switch configuration is loaded or not is controlled by the boot configuration. Loading of the routing switch configuration can be bypassed in the following ways:

- Use the Boot Monitor `flags` command and answer y when prompted:

  ```
  Do you want to use the factory default configuration (y/n)?
  ```

- Use the Run-Time CLI and issue this command:

  ```
  config sys set flags factory default true
  ```

When the configuration is bypassed, the routing switch boots in the factory default configuration except that the boot configuration settings have been loaded in stage 2. Bypassing the routing switch configuration does not affect the saved routing switch configuration; the configuration is simply not loaded.

Flash/PCMCIA File System

This section describes the flash/PCMCIA file system in an Accelar 1000 Series routing switch. The flash file system holds executable images and the switch configuration. The following sections are included:

- Flash memory organization (this page)
- Description of the file types (page 2-6)

Flash Memory Organization

There are two onboard flash memory devices on the Accelar routing switch: the Boot Flash and the System Flash. On Accelar 1200 Series switches, optional PCMCIA flash cards can be used. These devices are described in the following sections.

Boot Flash

The Boot Flash is 512 kilobytes (KB) and is divided into reserved areas for the boot monitor image and the routing switch configuration.

Boot Monitor Image

The boot monitor image is not directly user accessible. It is updated using a special boot monitor updater that writes to the area reserved for the boot image.

Switch Configuration (config and nvram)

The routing switch configuration is written whenever a save operation is performed on the configuration of the device. By default, the routing switch configuration is stored in a reserved area in the boot flash, although it is possible to specify alternative locations in the file system for the switch configuration.

- In the Boot Monitor CLI, use this command:

  ```
  choice
  ```

- In Run-Time CLI, use:

  ```
  config sys set config <choice>
  ```

The area reserved in boot flash for the switch configuration is accessed by the file system commands using the config or nvram file names. Both config and nvram names refer to the same file. Note that the switch configuration is read only when the run-time image loads.

System Flash (flash:)

The System Flash is 4 megabytes (MB) and is primarily used for run-time images, the system log, configuration files, and other general storage. The System Flash is divided into 64K blocks. Files stored in System Flash are always stored in an integral number of blocks.

Files stored in the System Flash are numbered sequentially starting with the numeral 1. Files can be assigned names by the user or referenced by their ordinal position in flash. The file naming convention for System Flash files is "flash:<filename>" or "flash:<file#>."

For example, flash:3 and flash:acc2_0_0 both refer to files in System Flash. In the first case, it is the third file in System Flash; in the latter case, it is the file named acc_2_2_0 in System Flash.

PCMCIA (pcmcia:)

Accelar 1200 Series routing switches can use an optional PCMCIA flash memory card. PCMCIA cards can be used for general storage for all file types and are a convenient way of moving files between switches because they are portable.

The PCMCIA card used in the Accelar 1200 and 1250 switches is the XLR1299PC PCMCIA Flash Memory Module. It has a capacity of 4 MB of memory with a block size of 128K.

As with System Flash, files stored on PCMCIA are numbered sequentially starting with 1 and can be given file names. The file naming convention for PCMCIA files is “pcmcia:<filename>” or “pcmcia:<file#>.”

File Types

Although System Flash and PCMCIA are primarily used for run-time images, configuration files, and the system log, they are also used to store other types of files. The following sections describe the various types of files that can be stored in the System Flash and PCMCIA. For a given file, the file type is reflected in the flags in a directory listing.

Executables

Executables are images that are executed by the Accelar 1000 Series CPU. The two most common executables needed by users are run-time images and boot monitor updaters. Note that executables can be stored in the flash file system in zipped format to conserve space. If necessary, the routing switch will automatically unzip the file upon execution.

Run-Time Images

The run-time image is an executable file that executes after the boot monitor image, initializing the I/O modules and providing full routing switch functionality. Run-time images can be stored and executed from System Flash and PCMCIA.

Boot Monitor Updaters

The boot monitor image is low-level code that initializes the devices on the Silicon Switch Fabric Module and starts the boot process. The boot monitor image is updated by executing a boot monitor updater that replaces the image stored in Boot Flash.

Log Files

Console information, warning, and error messages are logged to a log file. The log file is always stored in the System Flash. If no log file is present when the run-time image executes, a new log file is created. Log files are 128K, divided into two 64K banks. When the second bank fills, the first bank is erased and used again.

Configuration Files

In addition to the area reserved in Boot Flash for the switch configuration, configuration files can be stored and used in System Flash and PCMCIA.

Script Files

Script files are ASCII-based text files containing CLI commands that can be read by the switch and the commands executed as though they were typed at a console session.

Trace Logs

For debugging purposes, the routing switch creates a trace log with diagnostic messages. The trace log is not normally activated, so it is not normally accessed by end users. The file system commands refer to the reserved “trace” area for the trace log, so this information is presented for completeness.

Management Tools

You can use five management tools to monitor and manage your Accelar routing switch:

- Accelar Device Manager (this page)
- Accelar VLAN Manager (this page)
- Boot Monitor Command Line Interface (CLI) (page 2-9)
- Run-Time Command Line Interface (CLI) (page 2-9)
- Accelar Configuration Page (page 2-10)

Accelar Device Manager

Accelar Device Manager is an SNMP-based graphical user interface tool designed to manage single devices. In order to use Accelar Device Manager, you must have network connectivity to a management station running Accelar Device Manager on one of the supported platforms.

Accelar Device Manager is the most robust management tool in the Accelar Management Software suite and provides all the functionality you need to manage a single device, including the ability to create policy-based VLANs.

For more information about using Accelar Device Manager to configure the switch, refer to *Reference for Accelar Management Software Switching Operations* and *Reference for Accelar Management Software Routing Operations.*

Accelar VLAN Manager

Accelar VLAN Manager is an SNMP-based graphical user interface tool designed to manage VLANs across multiple devices. In order to use Accelar VLAN Manager, you must have network connectivity to a management station running Accelar VLAN Manager on one of the supported platforms.

Although you can manage VLANs through Accelar Device Manager, it is quicker and easier to manage multiple devices using Accelar VLAN Manager. When changes are made to your network configuration, STG IDs and VLAN IDs are automatically synchronized.

For information about using VLAN Manager, refer to Chapter 6 of *Reference for Accelar Management Software Switching Operations.*

Boot Monitor Command Line Interface (CLI)

The Boot Monitor CLI contains commands that enable you to configure boot options and manage files in flash memory. Changes that can be made and saved within the Boot Monitor CLI are boot choices, flags, IP configuration, and TFTP information. To access the Boot Monitor CLI, do *one* of the following:

- Interrupt the boot sequence by pressing a key when the following prompt is displayed:

  ```
  Press any key to stop autoboot.
  ```

- From the Run-Time CLI, enter the following commands, then reboot:

  ```
  Accelar-1100# config sys set flags autoboot false
  Accelar-1100# save
  ```

When you enter the Boot Monitor CLI, the following prompt is displayed:

```
monitor>
```

For information about the boot load process, refer to page 2-1. For boot monitor command information, refer to the section on the Boot Monitor CLI in *Reference for the Accelar 1000 Series Command Line Interface*.

Run-Time Command Line Interface (CLI)

The Run-Time CLI performs basic configuration tasks for SNMP management. To access the Run-Time CLI, you need a direct connection to the switch from a terminal or PC. Use a null-modem cable to connect the console port (DTE DB-9 male interface) to a DTE terminal or PC. Communication parameters are: 9600 bps, 8 data bits, no parity, 1 stop bit, with hardware flow control.

For pinout information about required cables, refer to Appendix A in *Using the Accelar 1200/1250 Routing Switch* or *Using the Accelar 1100/1150 Routing Switch*, or Appendix B in *Using the Accelar 1050/1051 Routing Switch*.

You also can access the Run-Time CLI through a Telnet or rlogin session. To open a Telnet session from Accelar Device Manager, click on the Telnet icon from the tool bar.

For more information about the CLI, refer to *Reference for the Accelar 1000 Series Command Line Interface*.

Accelar Configuration Page

The Accelar Configuration Page is a Web-based graphical user interface tool that operates in conjunction with a Web browser. It is designed to manage a single device and operates in most ways the same as Accelar Device Manager. However, it is not as complete in terms of functionality as Accelar Device Manager, and it is instead intended for use as a management tool to access and monitor devices on your network from various locations.

To access your switch, enter the device DNS name or IP address in the location field of a Web browser, such as Netscape Navigator. For example:

http://Accelar_1200 or *http://10.10.20.1/*

For more information about using the Accelar Configuration Page, refer to Chapter 10 of *Reference for Accelar Management Software Switching Operations*.

Note: To access a Web browser from Accelar Device Manager, click on the Web icon from the tool bar:

Accelar Access Levels and Passwords

The Accelar 1000 Series devices employ a security scheme with up to five levels of management access. The possible levels of security access are:

- Read-Only
- Layer 2 Read-Write functionality
- Layer 3 Read-Write functionality
- Read-Write for all functionality
- Read-Write-All

Read-Only Access

Read-Only access allows the manager to view the device settings, but changes are not allowed.

Layer 2 Read-Write Access

Layer 2 Read-Write access allows the manager to view and edit device settings dealing with layer 2 (bridging) functionality. The layer 3 settings (such as OSPF, DHCP) are not settable.

Layer 3 Read-Write Access

Layer 3 Read-Write access allows the manager to view and edit device settings dealing with layer 2 and layer 3 (routing) functionality.

Read-Write Access

Read-Write access allows the manager to view and edit most device settings. The only device settings that cannot be changed with Read-Write access are the security and password settings.

Read-Write-All Access

Read-Write-All access allows all the privileges of Read-Write access and the ability to change the security settings. The security settings include access passwords and the Web-based management user names and passwords.

CLI Access

When an Accelar 1000 Series routing switch is accessed for management, the user is prompted for a login and a password. The values for login and password for the console and Telnet sessions can be edited in the CLI or by using Device Manager.

- In the command line interface (CLI), set or reset CLI login and access passwords for the routing switch using this command:

  ```
  config cli password
  ```

- In Device Manager, from the main menu, choose Edit > Security > CLI to access the CLI window.

Note: You must have Read-Write-All privileges to view or change the CLI logins and passwords.

SNMP Community Strings

When opening a device using Accelar Device Manager, the user must specify both of the SNMP community strings. These community strings are used by the routing switch to authenticate received SNMP requests.The community strings specified determine the access level granted to the device.

In the CLI, set or reset SNMP community strings and access passwords for the routing switch using this command:

```
config sys set snmp community
```

In Device Manager, use the Edit Security SNMP window, accessed by selecting Edit > Security > SNMP, to set the community strings and access passwords.

Note: You must have Read-Write-All privileges to view or change SNMP community strings.

Web Parameters

In the command line interface (CLI), set or reset passwords for use with the Accelar Configuration Web Page using the following command:

```
config web-server set password <ro|rw|rwa> <username><passwd>
```

In Device Manager, use the Web window, accessed by choosing Edit > Security >Web, to configure Web passwords.

Note: You must have Read-Write-All privileges in order to change Web passwords.

Management Access

The Management Access feature allows you to control system access by setting access lists for services to prevent or allow access to the switch. You can specify which hosts or networks can access the switch through HTTP, rlogin, SNMP, and Telnet. If no access policies exist for a given service, the default is to allow the connection. Access policies apply to connections through the I/O port and not to the console or modem port. The lowest number indicates the highest precedence.

In the CLI, set these policies with these commands:

```
config ip access-policy
```

In Device Manager, set access policies through the Edit > Security > Access Policy window.

Port Lock

This feature allows a user to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. Locked ports cannot be modified in any way until the port is first unlocked.

In the CLI, this feature is set using these commands:

- `config ethernet <ports> lock <true|false>` to set individual ports or groups of ports.
- `config sys set portlock <on|off>` to globally turn port locking on or off.

In Device Manager, set port locking and unlocking through the Edit > Security > Port Lock window.

Chapter 3
Layer 2 Networking Concepts

This chapter discusses advanced layer 2 (switching) networking concepts with a special emphasis on how these concepts are implemented in an Accelar routing switch. This chapter assumes that you are familiar with the basics of bridging/switching, routing, and Spanning Tree Protocol. Topics covered in this chapter include:

- VLANs and VLAN types (this page)
- Port-based and policy-based VLANs, IP-subnet-based VLANs (starting on page 3-2)
- VLAN tagging and port types (page 3-8)
- Spanning Tree Protocol, spanning tree groups, and Accelar port FastStart (page 3-10)
- Traffic prioritization (page 3-12)
- Accelar 1000 Series VLAN specifics: special VLANs, defaults, and rules (page 3-14)
- Multi-Link Trunking (page 3-17)
- Network management (page 3-22)

VLANs

In a traditional shared-media network, traffic generated by a station is propagated to all other stations on the local segment. For a given station on shared Ethernet, the local segment is the “collision domain” because traffic on the segment has the potential to cause an Ethernet collision. The local segment is also the “broadcast domain” because any broadcast is sent to all stations on the local segment.

Ethernet bridges and switches divide a network into smaller “collision domains,” but they do not affect the broadcast domain. In simple terms, a virtual local area network (VLAN) can be thought of as a mechanism to fine-tune broadcast domains.

A VLAN is a collection of switch ports that make up a single broadcast domain. AVLAN can be defined for a single switch, or it can span multiple switches. VLANs are logical entities created in the software configuration to control traffic flow and ease the administration of moves, adds, and changes on the network.

On a given switch, a VLAN is one of two types: port-based or policy-based.

Port-Based VLANs

A port-based VLAN is a VLAN in which the ports are explicitly configured to be in the VLAN. When creating a port-based VLAN on a switch, you assign a VLAN identification number (VLAN ID) and specify which ports belong to the VLAN. The VLAN ID is used to coordinate VLANs across multiple switches. The mechanism for coordinating VLANs is described in the “VLAN Tagging and Port Types” section later in this chapter.

The example in Figure 3-1 shows two port-based VLANs: one for the marketing department and one for the sales department. Ports are assigned to each port-based VLAN. A change in the sales area can move the sales representative at port 3/1 (the first port in the I/O module in chassis slot 3) to the marketing department without moving cables. With a port-based VLAN, the network manager only needs to indicate in Accelar Device Manager that port 3/1 in the sales VLAN now is a member of the marketing VLAN.

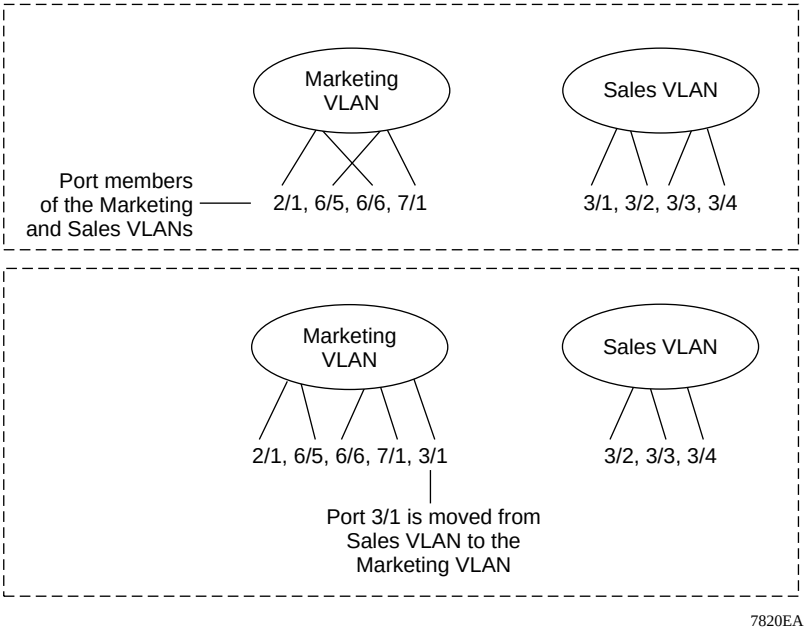
Marketing VLAN
Sales VLAN
Port members of the Marketing and Sales VLANs
2/1, 6/5, 6/6, 7/1
3/1, 3/2, 3/3, 3/4
Marketing VLAN
Sales VLAN
2/1, 6/5, 6/6, 7/1, 3/1
3/2, 3/3, 3/4
Port 3/1 is moved from Sales VLAN to the Marketing VLAN
7820EA

Figure 3-1. Example of Port-Based VLANs

Policy-Based VLANs

A policy-based VLAN is a VLAN in which ports are dynamically added to the VLAN based on the traffic coming into the port.

In a policy-based VLAN, ports are designated as always a member, never a member, or a potential member of the VLAN. When a port is designated as a potential member of the VLAN, the incoming traffic is monitored. When the incoming traffic matches the policy, the port is dynamically added to the VLAN. Potential member ports that have joined the VLAN are removed (“aged out”) from the VLAN if no traffic matching the policy is received within the aging time.

A port's membership in a VLAN is determined by the traffic coming into the port; therefore, Bay Networks recommends that at least some ports be designated as always a member of the VLAN. One situation in which a port should be designated always a member of a VLAN is if a server or router connects to the port. If a server is connected to a port that is only a potential member and the server sends out very little traffic, a client will fail to reach the server if the server port has timed out of the VLAN.

Accelar 1000 Series routing switches support policy-based VLANs based on the source MAC address, the network protocol, or the source IP subnet.

Note: A port can belong to multiple VLANs.

Source MAC-Based VLANs

As with all policy-based VLANs, using source MAC address VLANs allows the Accelar routing switch to associate frames with a VLAN based on the frame content. With source MAC-based VLANs, a frame is associated with a VLAN if the source MAC address is one of the MAC addresses explicitly associated with the VLAN by adding it to a list of MAC addresses that comprise the VLAN. However, because it is necessary to explicitly associate MAC addresses with a source MAC-based VLAN, the administrative overhead can be quite high.

Source MAC-based VLANs are used in situations where users want to enforce a MAC level security scheme to differentiate groups of users. For example, in a university environment, the students will be part of a student VLAN with certain services and access privileges, and the faculty will be part of a source MAC-based VLAN with faculty services and access privileges. Therefore, a student and a faculty member could plug into the same port but have the appropriate services. In order to provide the correct services throughout the campus, the source MAC-based VLAN would need to be defined on routing switches throughout the campus, which entails administrative overhead.

Note: ARU1 through ARU3 hardware does not support routing on a source MAC address-based VLAN.

Protocol-Based VLANs

As an example of using a protocol-based VLAN, a network manager can create a VLAN for the IPX protocol and place ports carrying substantial IPX traffic into this new VLAN. In Figure 3-2, the network manager has placed ports 7/1, 3/1, and 3/2 in an IPX VLAN. These ports still belong to their respective marketing and sales VLANs, but they are new members of the IPX VLAN also. This arrangement localizes traffic and ensures that only three ports will be flooded with IPX packets.

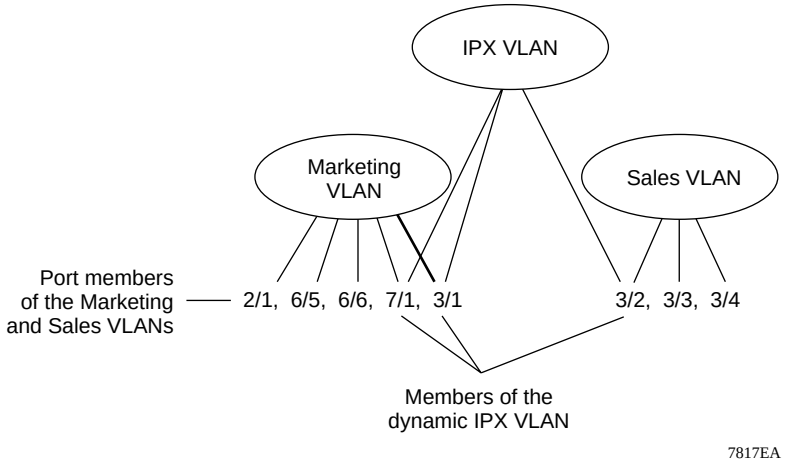

Figure 3-2. Example of a Dynamic VLAN Based on Protocol

The Accelar routing switch supports the following standard protocol-based VLANs:

- IP (ip)
- Novell IPX on Ethernet 802.3 frames (ipx802dot3)
- Novell IPX on IEEE 802.2 frames (ipx802dot2)
- Novell IPX on Ethernet SNAP frames (ipxSnap)
- Novell IPX on Ethernet Type 2 frames (ipxEthernet2)
- AppleTalk on Ethernet Type 2 and Ethernet Snap frames (AppleTalk)
- DEC LAT protocol (decLat)
- Other DEC protocols (decOther)
- IBM SNA on IEEE 802.2 frames (sna802dot2)

- IBM SNA on Ethernet Type 2 frames (snaEthernet2)
- NetBIOS Protocol (netBIOS)
- Xerox XNS (xns)
- Banyan VINES (vines)
- IP version 6 (ipv6)
- Reverse Address Resolution Protocol (RARP)

  RARP is a protocol used by some old diskless devices to obtain IP addresses by providing the MAC layer address. Creating a VLAN based on RARP allows controlling the RARP broadcast to the ports that would lead to the RARP server.

User-Defined Protocols

In addition to the standard protocols, user-defined protocol-based VLANs are supported. For user-defined protocol-based VLANs, the user specifies the Protocol Identifier (PID) for the VLAN. Any frames that match the specified PID in any of the following ways are assigned to that user-defined VLAN:

- The ethertype for Ethernet type 2 frames
- The PID in Ethernet Snap frames
- The DSAP or SSAP value in Ethernet 802.2 frames

The predefined policy-based PIDs are reserved and are not available for user-defined PIDs. Table 3-1 lists the reserved PIDs.

Table 3-1. Reserved PIDs for User-Defined Protocol-Based VLANs

PID (hex)	Comments
04xx, xx04	sna802dot2
F0xx, xxF0	netBIOS
0000-5DC	Overlaps with 802.3 frame length
0600, 0807	xns
0BAD	VINES
4242	IEEE 802.1D BPDUs
6000-6003, 6005-6009	decOther
6004	decLat

Table 3-1. Reserved PIDs for User-Defined Protocol-Based VLANs (continued)

PID (hex)	Comments
0800, 0806	ip
8035	RARP
8038	decOther
809B, 80F3	AppleTalk
8100	Reserved by IEEE 802.1Q for tagged frames
8137, 8138	ipxEthernet2 and ipxSnap
80D5	snaEthernet2
86DD	ipv6
8808	IEEE 802.3x pause frames
9000	Used by diagnostic loopback frames

Source IP Subnet-Based VLANs

Accelar switches with version -A or later I/O modules (ARU2 ASICs) also support policy-based VLANs based on IP subnets. Access ports can be assigned to multiple subnet-based VLANs. A frame's membership in a subnet-based VLAN is based on the IP source address associated with a mask. Subnet-based VLANs are optionally routable. Using source IP subnet-based VLANs, multiple workstations on a single port can belong to different subnets, similar to multinetting.

Note: IP subnet-based VLANs should not be used on segments that act as a transit network.

However, care should be exercised when using subnet-based VLANs. In the network example in Figure 3-3, when station 1 sends an IP frame to station 2, it will not arrive. Switch B will not assign this frame to either subnet VLAN 16 or 32 because of the IP source address 10.10.48.1.

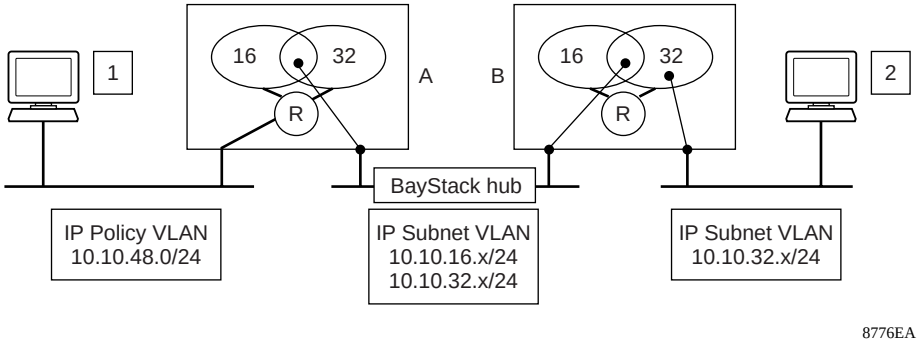

Figure 3-3. Example of IEEE 802.1Q Tagged Frame Format

VLAN Tagging and Port Types

Accelar 1000 Series routing switches support the IEEE 802.1Q specification for "tagging" frames. The specification defines a method for coordinating VLANs across multiple switches. In the specification, an additional 4-octet ("tag") header is inserted in a frame after the source address and before the frame type as shown in Figure 3-4. The tag contains the VLAN ID with which the frame is associated. By coordinating VLAN IDs across multiple switches, VLANs can be extended to multiple switches.

6 octets	6 octets	4 octets	2 octets	46-1500 octets	4 octets
Destination address	Source address	VLAN header: (VPID +VCI) TR-encap RESET*	Pkt type	Data	FCS

* If the source frame's data is in token ring format, and is required to be maintained in token ring format in transit across the VLAN, the TR-encap flag is set. If the source frame's data is not in token ring format, the TR-encap flag is reset.

7808EA

Figure 3-4. Example of Explicit Encapsulation Tagging

802.1Q Tagged Ports

Tagging a frame adds four octets to a frame, making it bigger than the traditional maximum frame size. Tagged frames that are bigger than the traditional maximum frame size are sometimes referred to as “baby giant” frames. If a device does not support IEEE 802.1Q tagging, it may have problems interpreting tagged frames and receiving baby giant frames.

In the Accelar routing switches, whether or not tagged frames are sent or received is configured at the port level. Tagging is set as true or false for the port and applied to all VLANs on that port.

An Accelar port with tagging enabled is a port from which all frames sent are tagged. Because all frames are explicitly tagged with a VLAN ID, tagged ports are typically used to multiplex traffic belonging to multiple VLANs to other IEEE-802.1Q-compliant devices. An Accelar tagged port can be configured to discard untagged frames or to associate them with a VLAN. In the latter case, when an untagged frame is received on a tagged port, it is sent to the user-specified VLAN.

An Accelar port with tagging disabled is a port that does not send tagged frames. A non-tagged port is used to connect Accelar routing switches to devices that do not support IEEE 802.1Q tagging. If a tagged frame is forwarded out a port with tagging set to false, the Accelar routing switch removes the tag from the frame before sending it out the port. A port with tagging set to false when receiving frames can be configured to discard tagged frames or to associate them with the VLAN specified in the tag.

Explicit Tagging Versus Implicit Tagging

When an Accelar routing switch receives a frame, how the frame is forwarded is based on the VLAN on which the frame is received and based on the forwarding options available for the VLAN. The frame is associated with a VLAN through either explicit or implicit tagging.

A frame is explicitly tagged if it is received on a tagged port and is tagged. In this instance, the frame is already associated with a VLAN in its tag.

A frame is implicitly tagged if the frame is received without a tag. The Accelar routing switch associates the frame with a VLAN based upon the data content of the frame or the receiving port. Because no VLAN tag is present, VLAN membership is implied from the content of the frame itself. If you choose not to discard untagged frames on a tagged port, you must specify a port-based VLAN on STG1 as the default and the tagged port must be a member of that VLAN.

Accelar routing switches try to associate a frame with the source MAC address (source MAC-based VLAN), source IP address (source IP-subnet based VLAN), protocol-based VLANs, and then port-based VLANs. Untagged frames are associated with a VLAN according to the following criteria:

- Does the frame belong to a source MAC-based VLAN?
- Does the frame belong to a source IP-subnet VLAN?
- Does the frame belong to a protocol-based VLAN?
- What is the port-based VLAN of the receiving port?

IP Routing and VLANs

The Accelar routing switch supports IP routing on the following types of VLANs only:

- Port-based VLANs
- Source-IP subnet-based VLANs
- IP protocol-based VLANs

IP routing is not supported on source MAC-based VLANs or VLANs based on other protocols, including IP version 6 and user-defined protocol-based VLANs.

Spanning Tree Protocol, Groups, and FastStart

Path redundancy for VLANs is controlled by implementing the Spanning Tree Protocol (STP). A network may include multiple instances of STP. The collection of ports in one spanning tree instance is called a spanning tree group (STG). The Accelar routing switch supports Spanning Tree Protocol and multiple spanning tree instances, thus multiple spanning tree groups.

Spanning Tree Protocol

As defined in the IEEE 802.1D standard, the Spanning Tree Protocol detects and eliminates logical loops in a bridged or switched network. When multiple paths exist, the spanning tree algorithm configures the network so that a bridge or switch uses only the most efficient path. If that path fails, the protocol automatically reconfigures the network to make another path become active, thus sustaining network operations.

Accelar Spanning Tree Groups

Accelar 1000 Series routing switches support the Spanning Tree Protocol. In addition, a routing switch can support multiple spanning tree groups within the same box; that is, the routing switch can participate in the negotiation for multiple spanning trees. Figure 3-5 shows multiple spanning tree groups.

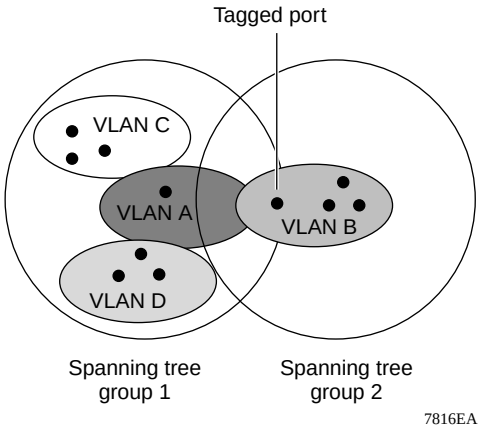

Figure 3-5. Multiple Spanning Tree Groups

The ports associated with a VLAN must be contained within a single spanning tree group. Not allowing a VLAN to span multiple spanning tree groups avoids problems with spanning tree blocking ports and causing a loss of connectivity within a VLAN.

Each untagged port can belong to one and only one spanning tree group. The 802.1Q tagged ports can belong to more than one spanning tree group. When a tagged port belongs to more than one spanning tree group, the spanning tree Bridge Protocol Data Units (BPDUs) are sent as tagged frames with a VLAN ID. Because tagged BPDUs are not a part of the 802.1D standard, not all devices can interpret tagged BPDUs.

Accelar Spanning Tree FastStart

Spanning Tree FastStart is an enhanced port mode supported by Accelar 1000 Series routing switches. If Spanning Tree FastStart is enabled on a port, the port is brought up more quickly following the routing switch initialization or a spanning tree change. The port goes through the normal blocking and learning states before the forwarding state, but the hold times for these states is the bridge hello timer (2 seconds by default) instead of the bridge forward delay timer (15 seconds by default). Enabling FastStart allows for faster convergence upon topology change. FastStart is useful on access ports where there may be only one device connected to the switch (as in workstations with no other spanning tree devices), and it may not be desirable to wait for the usual 30 to 35 seconds for spanning tree initialization and bridge learning.

Note: Use Accelar Spanning Tree FastStart with caution. A loop condition may exist until a Bridge Protocol Data Unit is seen on the port configured for FastStart. This procedure is contrary to that specified in the IEEE 802.1D standard for Spanning Tree Protocol (STP), in which a port enters blocking state following the initialization of the bridging device or from the disabled state when the port is enabled through configuration.

Traffic Prioritization

Accelar 1000 Series routing switches prioritize traffic using queues and headers. As each packet is forwarded through the switch fabric, a header is attached. The header contains prioritization information set by the forwarding engine on the ingress port when the packet is received. Each time a packet is forwarded within the switch, it is placed in either a high-priority or low-priority queue depending upon the priority information in the internal packet header. At each stage within the switch, packets in high-priority queues are sent before packets in low-priority queues.

Traffic prioritization allows information technology managers to prioritize mission-critical transmissions. With Accelar routing switches, you can set traffic prioritization to assign a packet or data stream a high-priority queue so that it goes through the network with minimal latency. The Accelar routing switch has four queues for traffic:

- Unicast low
- Unicast high
- Multicast low
- Multicast high

Traffic priority is especially critical for multimedia—video in particular. Using Accelar Device Manager or Accelar VLAN Manager, you can prioritize packets to provide more bandwidth for traffic that requires it. For example, you could assign a higher priority to use more bandwidth for voice and video multimedia traffic. In this way, you can control the delivery of multimedia traffic to eliminate jerky transitions. The result of setting a higher priority for multimedia traffic is a smoother image and better sound quality.

Setting Priority

An Accelar switch can operate in either of two modes of traffic priority: Best Effort mode or Priority mode. The factory default setting is Best Effort mode. The following differences exist between these modes:

- In Best Effort mode, all traffic is treated with the same priority.
- In Priority mode, high-priority traffic flows through the switch fabric using a high-priority data path. Output buffers are reserved for high-priority traffic.

High-priority traffic can be enabled based on a per-port, per-MAC, per-VLAN, or per-flow basis as follows:

- When a port is set to high-priority mode, all traffic received on this port is assigned a high priority.
- When a MAC address is set to high-priority mode, all traffic from the MAC address is assigned a high priority.
- When a VLAN is set to high-priority mode, frames received on any of the active ports of the VLAN are assigned a high priority.

- An IP flow record can be used to assign high switching priority to an IP packet based on its source and destination IP addresses, protocol type source port number, and destination port number.

When a high-priority frame is sent out a tagged port, the 3-bit User Priority field in the IEEE VLAN tag is set to 7. A normal priority frame has a User Priority of 0. Any received tagged frames with a User Priority greater than 2 are treated as high priority.

The Accelar 1000 Series routing switch is also compliant to the IEEE 802.1p standard. This standard specifies a priority bit on 802.1Q tagged frames. Upon receipt of these frames, the Accelar switch will place frames with a priority of 2 or greater into the high-priority queue.

For information about layer 3 prioritization, refer to Chapter 6, “IP Filtering.”

Accelar 1000 Series VLANs

This section describes specifics of how VLANs are implemented in Accelar 1000 Series routing switches. In particular, this section describes two special, predefined VLANs in Accelar routing. It also summarizes the defaults and rules regarding VLAN creation on Accelar 1000 Series routing switches.

VLAN Rules

Accelar 1000 Series routing switch VLANs operate under the following basic set of rules:

- Accelar routing switches support 123 VLANs in addition to the default VLAN. VLAN IDs range in value from 1 to 4094.
- For every STG group that you create, you reduce by one the number of VLANs that you can create (up to 123). For example, the maximum number of VLANs supported is 123; so if you create 12 STGs, you would then be allowed 111 user-defined VLANs.
- For every IGMP-snoop group you create, you also reduce by one the number of VLANs that you can create.

Note: IGMPv1 snooping requires hardware with ARU2 or above ASICs, IGMPv2 snooping requires hardware with ARU3 ASICs. If all hardware modules in a switch have ARU3 ASICs and the switch is running release 2.0 or later, it is no longer true that the total number of available VLANs is reduced by one for each multicast group per VLAN. Instead, up to 1024 source IP subnet/multicast group combinations per switch are allowed.

- For every VLAN with Multi-Link Trunking that you create, you reduce by four the number of available VLANs.
- A VLAN cannot span multiple spanning tree groups; that is, the ports in the VLAN must all be within one spanning tree group. Spanning Tree Group IDs can range in value from 1 to 128.
- An untagged port can belong to one and only one port-based VLAN. A port in a port-based VLAN can belong to other policy-based VLANs.
- An untagged port can belong to one and only one policy-based VLAN for a given protocol. For example, a port can belong to only one policy-based VLAN where the policy is IPX802dot2 protocol.
- A frame's membership in an IP subnet-based VLAN takes precedence over the protocol-based VLAN; the protocol-based VLAN takes precedence over the port-based VLAN.
- The IP subnet-based VLAN should not be assigned to a transit network, a network bridging two other subnets.
- Tagged ports can belong to multiple VLANs and multiple spanning tree groups. When a tagged port belongs to multiple spanning tree groups, the BPDUs are tagged for all spanning tree groups except for the default spanning tree group. Under the default configuration, the default spanning tree group is number 1.

Special VLANs

Accelar 1000 Series routing switches have two predefined VLANs that behave differently from user-defined VLANs. These VLANs are the default VLAN and the unassigned VLAN.

Default VLAN

Accelar 1000 Series routing switches are factory configured with all ports in a port-based VLAN called the default VLAN. With all ports in the default VLAN, the switch behaves like a layer 2 switch.

The VLAN ID of the default VLAN is always 1, and it is always a port-based VLAN. The default VLAN cannot be deleted. When a user-defined port-based VLAN is deleted, all ports are moved back into the default VLAN to help maintain connectivity.

Unassigned VLAN

Internally, Accelar routing switches support a placeholder for ports that is called an unassigned port-based VLAN. This unassigned concept is used for ports that are removed from all port-based VLANs. Ports can belong to policy-based VLANs as well as to the unassigned VLAN. If a frame does not meet any policy criteria and there is no underlying port-based VLAN, the port belongs to the unassigned VLAN and the frame is dropped. Ports only in the unassigned VLAN have no spanning tree group association, so these ports do not participate in Spanning Tree Protocol negotiation; that is, no BPDUs are sent out of ports in the unassigned VLAN. Isolated routing ports are an example of ports in an unassigned VLAN because they are not associated with a VLAN and they do not participate in spanning tree negotiations.

Because it is an internal construct, the unassigned VLAN cannot be deleted. If a user-defined spanning tree group is deleted, the ports are moved to the unassigned VLAN and can later be assigned to another spanning tree group. Moving the ports to the unassigned VLAN avoids creating unwanted loops and duplicate connections. If routing is disabled in these ports, the port is completely isolated and no layer 2 or layer 3 functionality is provided.

The concept of unassigned VLANs is useful for security concerns or when using a port for monitoring a mirrored port.

Brouter Ports

Another special VLAN supported by the Accelar switch is a brouter port, which is actually a one-port VLAN. The difference between a brouter port and a standard IP protocol-based VLAN configured to do routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. Brouter ports are discussed in more detail in the section titled “Brouter Ports” on page 4-5.

Default Configuration

When you boot your Accelar 1000 Series routing switch, it will contain the following default configuration:

- A single, port-based VLAN is configured. The default VLAN has a VLAN identification number of 1 and is bound to the default spanning tree group.
- All ports are in a single spanning tree group. The spanning tree group number is 1. The default spanning tree group is 802.1D compliant, and its BPDUs are never tagged.
- Spanning Tree FastStart is disabled on all ports.
- No interfaces in the default configuration are assigned IP addresses.
- Traffic priority for all ports is set to normal priority.
- All ports are nontagged ports.

Multi-Link Trunking

Multi-Link Trunking (MLT) is a point-to-point connection that aggregates multiple ports so they logically act like a single port with the aggregated bandwidth. Grouping multiple ports into a logical link allows the user to achieve higher aggregate throughput on a switch-to-switch or switch-to-server application. Multi-Link Trunking provides media and module redundancy.

A number of Bay Networks products implement MLT and will have different features and requirements based on the architecture of the device. For the Accelar routing switches, MLT has the following general features and requirements:

- MLT is supported on version -A hardware modules and chassis and above only.
- A chassis can have up to eight MLT connections.
- As many as four same-type ports can belong to a single Multi-Link Trunk (MLT).
- The ports in an MLT can span modules, providing module redundancy.
- MLT is supported on 10BASE-T, 100BASE-TX, 100BASE-FX, and Gigabit Ethernet ports.
- All ports in an MLT must be of the same media type (copper or fiber) and have the same settings (speed and duplex).

- All ports in an MLT must be in the same spanning tree group.
- MLT is compatible with the Spanning Tree Protocol.
- IEEE 802.1Q tagging is supported on an MLT.
- For bridge traffic, the algorithm that distributes traffic across an MLT is based on the source and destination MAC addresses.
- For routed traffic, the algorithm that distributes traffic across an MLT is based on the source and destination IP addresses.

Keep in mind that setting up MLTs reduces the number of VLANs available on the switch. An Accelar switch starts with 123 available VLANs. Every VLAN that includes one or more MLT reduces the total number of available VLANs by four.

Multi-Link Trunking Examples

Multi-Link Trunks allow you to group up to four switch ports together to form a link to another switch or server, thus increasing aggregate throughput of the interconnection between the devices (up to 8 Gb/s in full-duplex mode). Accelar switches can be configured with up to eight Multi-Link Trunks. When Spanning Tree Protocol is enabled, Multi-Link Trunking software detects misconfigured (or broken) trunk links and redirects traffic on the misconfigured or broken trunk link to other trunk members within that trunk. Figure 3-6 shows two trunks (T1 and T2) connecting switch S1 to switches S2 and S3.

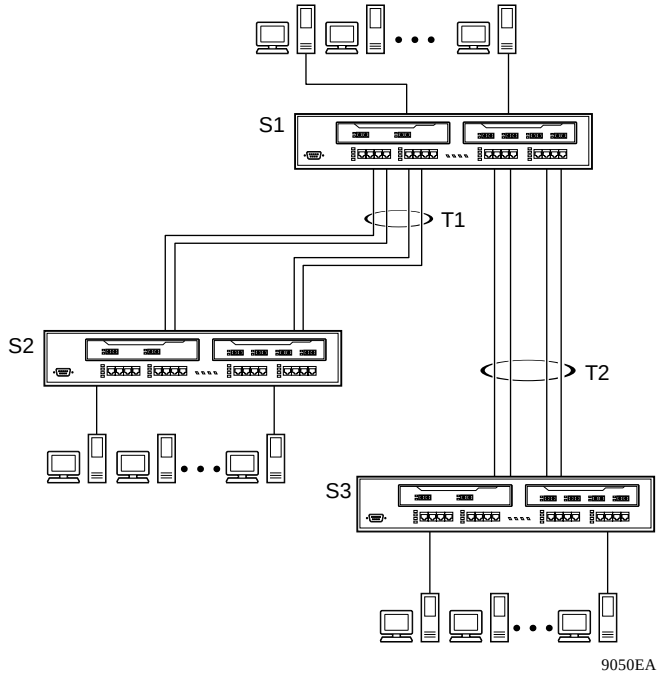

Figure 3-6. Switch-to-Switch MLT Configuration Example

Each of the trunks shown in Figure 3-6 can be configured with up to four switch ports to increase bandwidth. When traffic between switch-to-switch connections approaches single port bandwidth limitations, creating a Multi-Link Trunk can supply the additional bandwidth required to improve the performance.

Figure 3-7 shows a typical switch-to-server trunk configuration. In this example, file server FS1 utilizes dual MAC addresses, using one MAC address for each network interface controller (NIC). FS2 is a single MAC server (with a 4-port NIC) and is set up as trunk configuration T1.

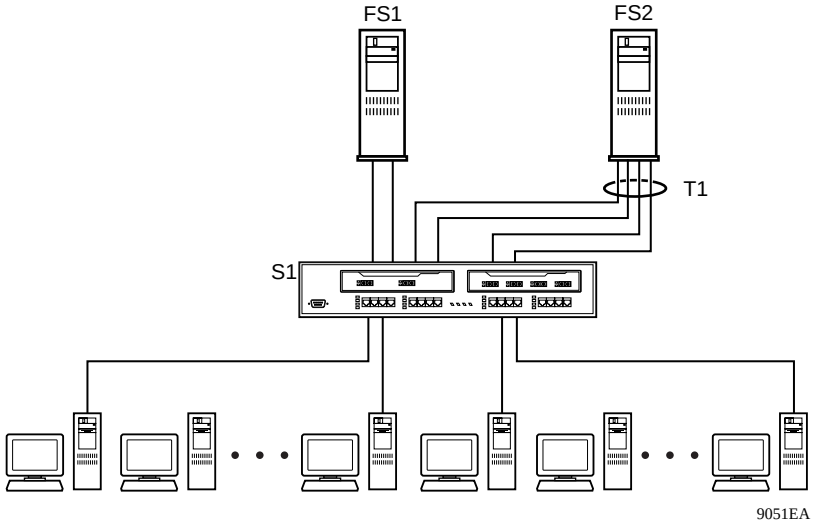

Figure 3-7. Switch-to-Server MLT Configuration Example

Client/Server Configuration Utilizing Multi-Link Trunks

Figure 3-8 shows an example of how Multi-Link Trunking can be used in a client/server configuration. In this example, both servers are connected directly to switch S1. FS2 is connected through a trunk configuration (T1). The switch-to-switch connections are through trunks (T2, T3, T4, and T5). Clients accessing data from the servers (FS1 and FS2) are provided with maximized bandwidth through trunks T1, T2, T3, T4, and T5. Trunk members (the ports making up each trunk) do not have to be consecutive switch ports; they can be selected randomly, as shown by T5.

With spanning tree *enabled* and trunks T2 and T3 in the same spanning tree group, one of the trunks (T2 or T3) acts as a redundant (backup) trunk to switch S2. With spanning tree *disabled*, trunks T2 and T3 must be configured into separate VLANs for this configuration to function properly.

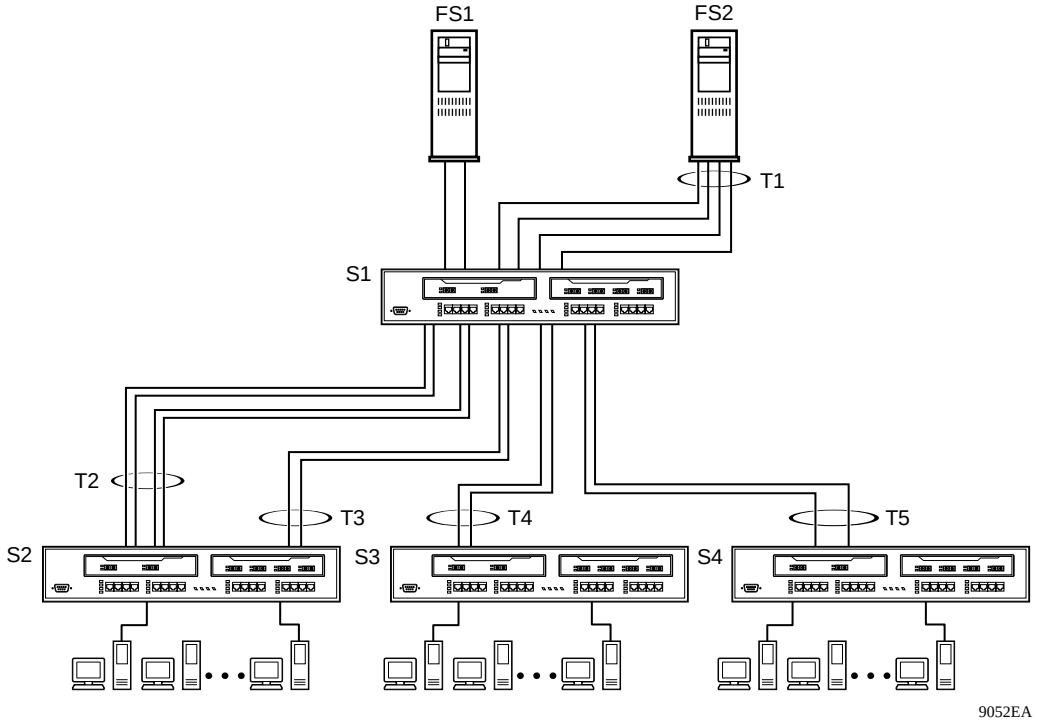

Figure 3-8. Client/Server Configuration Example

Ports that belong to the same Multi-Link Trunk operate as follows. All ports in the MLT must belong to the same spanning tree group if spanning tree is enabled. Identical Bridge Protocol Data Units (BPDUs) are sent out each port. The MLT port ID is the ID of the lowest numbered port. If identical BPDUs are received on all ports, the MLT mode is forwarding. If no BPDU is received on a port or if BPDU tagging and port tagging do not match, the individual port is taken offline. Path cost is inversely proportional to the active MLT bandwidth.

Network Management and Diagnostics

You can manage your network from two graphical user interfaces (GUIs) called the Accelar Device Manager and the Accelar VLAN Manager, from the Web, or from a command line.

- Accelar Device Manager is SNMP-based and runs on UNIX (Solaris, HP-UX, and IBM AIX), Windows® 95, Windows 98, and Windows NT® platforms.
- Accelar Device Manager is used to manage one device at a time, and Accelar VLAN Manager is used to manage VLANs across multiple devices at the same time.
- From a Web browser, you can manage Accelar routing switches using the Accelar Configuration Page. Using a Web browser, such as Netscape Navigator, you can enter the DNS name or IP address of your switch in the location field of the Web browser and bring up a management menu and graphical representation of your switch.

  For more information about using the Accelar Configuration Page, refer to *Reference for Accelar Management Software Switching Operations*.
- A command line interface (CLI) is available to perform tasks outside of the Device Manager or VLAN Manager graphical user interfaces. It is accessible via a console or any Ethernet port using Telnet or rlogin.

  For more information about the CLI, refer to *Reference for the Accelar 1000 Series Command Line Interface* (Bay Networks part number 202086-A).

With these management tools, you can view the trap log, use RMON to create and manage alarms, and use other diagnostic tools such as port mirroring to analyze traffic on a per-port basis and Syslog to map informational messages and warnings.

RMON

Remote monitoring (RMON) is a management information base (MIB) or a group of “management objects” that you use to “get” or “set” values using Simple Network Management Protocol (SNMP). Using the CLI or Accelar Device Manager, you can enable RMON globally; using Device Manager, you can also enable RMON on a port-by-port basis.

RMON has three major functions:

- Setting alarms for user-defined events
- Gathering real-time and historical Ethernet statistics
- Logging events
- Sending traps for events

The Accelar implementation of RMON lets you set alarms relating to specific events, or variables, that you select from a drop-down menu. You specify events associated with alarms to be set to either trap or log-and-trap. In turn, these alarms, when tripped, are trapped or logged.

Although all information is viewable from Accelar Device Manager, you can use any management application that supports SNMP traps residing on another device (such as HP OpenView running on a Sun workstation) to view RMON trap information remotely.

Port Mirroring

Accelar routing switches support the port mirroring management feature to analyze traffic.

Using port mirroring, you can specify a destination port on which you want to see mirrored traffic and specify the source ports from which traffic is mirrored. Any packets ingressing or egressing the specified ports are forwarded normally, and a copy of the packet is sent out the mirror port. You observe packet traffic at the destination port using a network analyzer—a *copy* of the packets can be captured and analyzed. Unlike with other methods used to analyze packet traffic, packet traffic is uninterrupted and packets flow normally through the destination port.

An Accelar 1000 Series routing switch can support mirroring for only two ports. When this feature is active, all packets received or transmitted on the port(s) specified by MirroredPortOne and/or MirroredPortTwo are copied to MirrorPort. The mirroring operation is nonintrusive.

In addition, the port mirroring feature can be used to monitor traffic for MAC addresses where traffic with a given source or destination MAC address is copied to the mirror port. So as not to see unintended traffic, remove the port to which you are mirroring from all VLANs; that is, move it to the unassigned VLAN.

Syslog

On any UNIX-based management platform, you can use the Syslog messaging feature of the Accelar routing switch to manage routing switch event messages. The Accelar syslog software supports this functionality by communicating with a counterpart software component named *syslogd* on your management workstation. The UNIX daemon syslogd is a software component that receives and locally logs, displays, prints, and/or forwards messages that originate from sources internal and external to the workstation. For example, syslogd on a UNIX workstation concurrently handles messages received from applications running on the workstation, as well as messages received from Accelar routing switches running in a network accessible to the workstation.

At a remote UNIX management workstation, syslogd does the following:

- Receives syslog messages from the Accelar routing switch
- Examines the severity code in each message
- Uses the severity code to determine appropriate system handling for each message
- Based on the severity code in each message, dispatches each message to any or all of the following destinations:
  - — Workstation display
  - — Local log file
  - — Designated printer
  - — One or more remote hosts

Internally the Accelar routing switch has four severity levels for log messages:

- Info
- Warning
- Error
- Fatal

Syslog supports eight different severity levels:

- Debug
- Info
- Notice
- Warning
- Error
- Critical
- Alert
- Emergency

Chapter 4
IP Interfaces and Router Management

The Accelar 1000 Series routing switches support wire-speed IP routing of frames. The routing protocols supported are RIP version 1 (RFC1058), RIP version 2 (RFC 1723), and OSPF version 2 (RFC 1583).

The router management features covered in this chapter apply regardless of which routing protocols are used and include router IP configuration, IP route table management, ARP configuration, ARP table management, BootP/DHCP relay configuration, and VRRP configuration.

The following topics are covered:

- IP addresses (this page)
- Types of IP routing (page 4-3)
- Static routes (page 4-6)
- Router management (page 4-6)
- Dynamic IP routing protocols (page 4-15)

IP Addresses

An IP address consists of 32 bits that have the form *network.host*. The *network* portion is a network number ranging from 8 to 24 bits. The *host* portion is the remaining 8 to 24 bits identifying a specific host on the network. The Internet Network Information Center (NIC) assigns the network portion of the IP address. Your network administrator assigns the host portion.

You specify IP addresses in dotted-decimal notation. To express an IP address in dotted-decimal notation, you convert each 8-bit octet of the IP address to a decimal number and separate the numbers by decimal points.

For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

Subnet Addressing

The concept of subnetworks (or subnets) extends the IP addressing scheme. Subnets are two or more physical networks that share a common network-identification field (the NIC-assigned network portion of the 32-bit IP address). Subnets allow you to further divide a network into multiple routed segments.

With subnets, you partition the host portion of an IP address into a subnet number and a "real" host number on that subnet. The IP address is then defined by *network.subnet.host*. Routers outside the network do not interpret the subnet and host portions of the IP address separately.

Routers inside a network containing subnets use a 32-bit subnet mask that identifies the extension bits. In *network.subnet.host*, the *subnet.host* portion (or the local portion) contains an arbitrary number of bits. The network administrator allocates bits within the local portion to subnet and host and then assigns values to subnet and host.

Supernet Addressing

A supernet is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the following block of contiguous 32-bit addresses (192.32.0.0 to 192.32.7.0 in dotted-decimal notation):

11000000 00100000 00000000 00000000
11000000 00100000 00000001 00000000
11000000 00100000 00000010 00000000
11000000 00100000 00000011 00000000
11000000 00100000 00000100 00000000
11000000 00100000 00000111 00000000
11000000 00100000 00000101 00000000
11000000 00100000 00000110 00000000

The supernet address is also referred to as the classless interdomain routing (CIDR) address. The supernet address for this block is 11000000 00100000 00000, the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an *address*/*mask* pair:

- *address* is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- *mask* is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/21.

Types of IP Routing

There are three types of router interfaces: physical router interfaces (also called isolated routing ports), virtual router interfaces (routing between VLANs), and brouter ports (routing and bridging on the same port). In an isolated routing port, an IP address is associated with a physical port. When routing on a VLAN, an IP address is assigned to the VLAN and is not associated with any particular physical port. This difference is an important distinction between isolated routing ports and routing between VLANs. Brouter ports are actually one-port VLANs that route IP packets and bridge nonroutable traffic.

Isolated Routing Ports

Any port in the Accelar 1000 Series routing switch can be configured as an isolated IP routing port as shown in Figure 4-1. In this mode, the port behaves like a traditional router port. The port only routes IP packets and does not perform any bridging. The IP address for the isolated routing port is associated with the physical port.

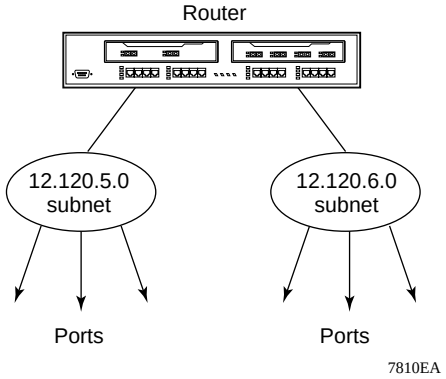

Figure 4-1. Routing Between IP Destination Addresses

Ports connected to the wide area routers or in the network backbones where there is no requirement for bridging non-IP traffic are most likely to be configured as isolated routing ports. If bridging of other protocols is required, you can configure a VLAN on a port or a set of ports and enable routing for that VLAN.

For SNMP or Telnet management, you can use any isolated router port interface address to access the routing switch.

Virtual Routing Between VLANs

Accelar routing switches support wire-speed IP routing between VLANs as shown in Figure 4-2. When routing is configured on a VLAN, an IP address is assigned to the VLAN that acts like a “virtual router interface” address for the VLAN. It is a virtual router interface in that it does not have an association with any particular port. The IP address can be reached through any of the ports in the VLAN, and it is the IP address for the gateway through which a frame is routed out of the VLAN. Routed traffic can be forwarded to another VLAN within the routing switch or to an isolated routing port.

When spanning tree is enabled in a VLAN, spanning tree convergence must have stabilized before the router protocol can begin. This requirement can lead to an additional delay in the forwarding of IP traffic.

Because a given port can belong to multiple VLANs (some of which are configured for routing on the switch and some of which are not), there is no longer a one-to-one correspondence between the physical port and the router interface.

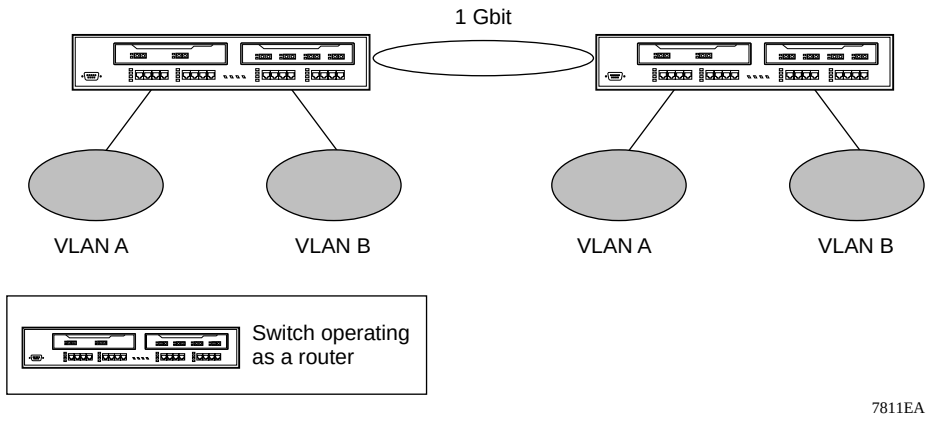

Figure 4-2. IP Routing Between VLANs

Virtual router interface addresses are also used for device management. For SNMP or Telnet management, you can use any virtual router interface address to access the routing switch as long as routing is enabled on the VLAN.

Brouter Ports

The Accelar switch also supports the concept of brouter ports. A brouter port is a single-port VLAN that differs from an isolated routing port in that it can route IP packets as well as bridge all nonroutable traffic. The difference between a brouter port and a standard IP protocol-based VLAN configured to do routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic and still be able to route the IP traffic. This feature removes any delays caused by spanning tree in routed traffic.

To create a brouter port, configure a routed IP policy-based single-port VLAN into spanning tree group 0. Subsequent VLANs on the port are part of the same spanning tree group. When spanning tree blocks the port, the IP policy-based VLAN will continue to forward (route) packets.

A brouter port is actually a one-port VLAN; therefore, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

Static Routes

Static routes are used to provide a mechanism to create routes to the destination IP address prefixes manually.

A static default route is used to specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This route is by definition a route with the prefix length of zero [RFC1812]. The routing switches can be configured with the default route statically, or they can learn it via a dynamic routing protocol.

Note: To create a default static route, the destination address and subnet mask must be set to 0.0.0.0.

Static routes can also be configured with a next hop that is not directly connected.

Router Management

The following sections describe various protocols used in router management. Topics include:

- Address Resolution Protocol (ARP) (this page)
- BootP/DHCP relay (page 4-8)
- UDP broadcast forwarding (page 4-11)
- Reverse Address Resolution Protocol (RARP) (page 4-12)
- Virtual Router Redundancy Protocol (VRRP) (page 4-12)

Address Resolution Protocol (ARP)

An IP router needs both a physical address and an IP address to transmit a datagram. In situations where the router knows only the network host's IP address, the Address Resolution Protocol (ARP) enables the router to determine a network host's physical address by binding a 32-bit IP address to a 48-bit MAC address. A router can use ARP across a single network only, and the network hardware must support physical broadcasts.

If a router wants to send a packet to a host but knows only the host's IP address, the router uses ARP to determine the host's physical address as follows:

1. The router broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.
2. All network hosts receive the broadcast request.
3. Only the specified host responds with its hardware address.

The router then maps the host's IP address to its physical address and saves the results in an address-resolution cache for future use. The router's ARP table displays the known MAC address to IP address associations. Static ARP entries can be created, and individual ARP entries can be deleted.

Using Proxy ARP

Proxy ARP allows a router to respond to an ARP request from a locally attached host or end station for a remote destination. It does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet on which the ARP request was received. The reply is generated only if the switch has an active route to the destination network.

Figure 4-3 is an example of proxy ARP operation. Host B could send an ARP request for Host C. The Accelar routing switch would respond to the ARP request with Host C's IP address but with its own MAC address.

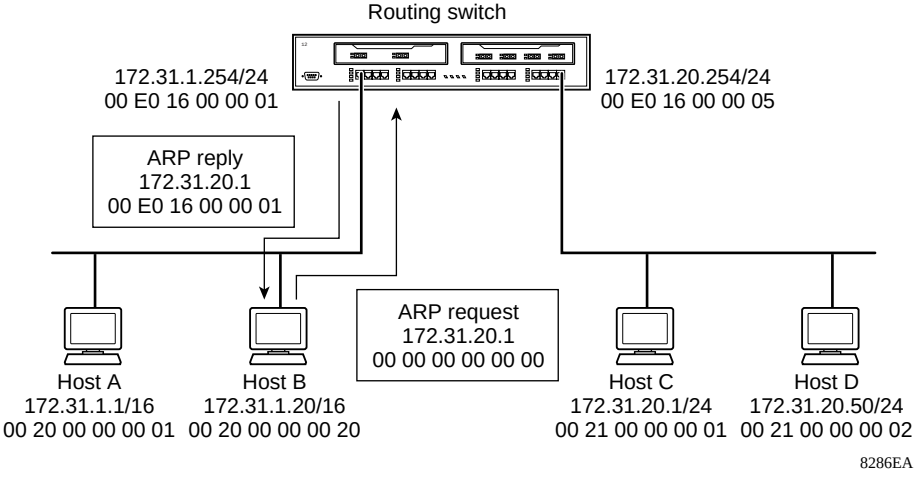

Figure 4-3. Proxy ARP Operation

Flushing Router Tables

For administrative and/or troubleshooting purposes, it is sometimes necessary to flush the routing tables. Accelar Device Manager provides facilities for doing this in two contexts: by VLAN and by port.

In a VLAN context, all entries associated with the VLAN will be flushed. In a port context, all entries associated with the port will be flushed.

BootP/DHCP Relay

Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), is used to dynamically provide host configuration information to the workstations. To lower administrative overhead, network managers prefer to configure a small number of DHCP servers in a central location. Using few DHCP servers requires the routers connecting to the subnets or VLANs/bridge domains to support the BootP/DHCP relay function so that hosts can get the configuration information from servers several router hops away.

Differences Between DHCP and BootP

The following differences between DHCP and BootP are specified in RFC 2131 and include functions that BootP does not address:

- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters needed to operate.

DHCP uses the BootP message format defined in RFC 951. A packet is classified as DHCP if the first four octets in the options field are 99, 130, 83, 99, and the fifth octet is 53. The first four octets are referred to as the “Magic Cookie”; the fifth is the DHCP message type code. The remainder of the options field consists of a list of tagged parameters that are called “options” (RFC2131).

Summary of DHCP Relay Operation

BootP/DHCP clients (workstations) generally use UDP/IP broadcasts to determine their IP addresses and configuration information. If such a host is on a network or a subnet segment (or VLAN) that does not include a DHCP server, the UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. The Accelar routing switches can be configured to

overcome this issue by forwarding the broadcasts to the server through isolated or virtual router interfaces. The router interfaces can be configured to forward DHCP broadcasts to other locally connected network segments or directly to the server's IP address. DHCP must be enabled on a per-routable-interface basis.

In Figure 4-4, an end station is connected to subnet 1, corresponding to VLAN 1. The Accelar routing switch connects two subnets via the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255) with the DHCP relay function configured, the Accelar routing switch forwards DHCP requests to subnet 2 or to the host address of the DHCP server, depending on the configuration.

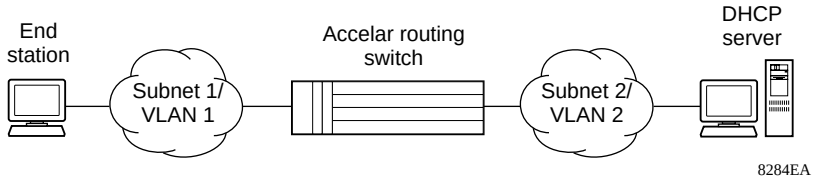

Figure 4-4. Example of DHCP Operation

Forwarding DHCP Packets

In the example shown in Figure 4-5, the *agent address* is: 10.10.1.2.

- To configure the Accelar routing switch to forward DHCP packets from the end station to the server, use 10.10.2.1. as the *server address*.

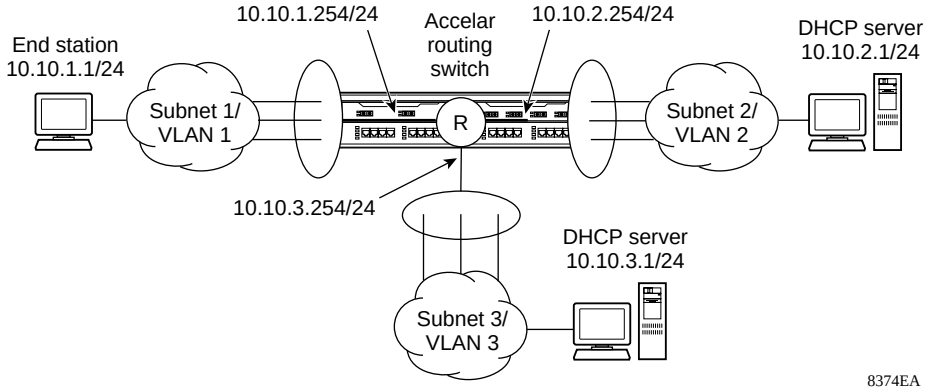

Figure 4-5. Forwarding DHCP Packets

All BootP broadcast packets, including DHCP packets that appear on the VLAN 1 router interface (10.10.1.2), will be forwarded to the DHCP server. In this case, the DHCP packets will be forwarded as unicast to the DHCP server's IP address.

- To forward BootP/DHCP packets as broadcast packets to VLAN 2, specify the IP address of the switch VLAN2 router interface (10.10.2.2.) as the server address.

Multiple BootP/DHCP Servers

Most enterprise networks use multiple BootP/DHCP servers for fault tolerance. The Accelar routing switches allow configuring to forward the BootP/DHCP requests to multiple servers. Up to 10 servers can be configured to receive copies of the forwarded/relayed BootP/DHCP messages.

If a DHCP client is connected to a routable interface, to configure DHCP requests to be sent to 10 different routable interfaces or 10 different server IP addresses, enable DHCP on the client (agent address) and then enable DHCP from the client to each of the interfaces or IP addresses (server addresses).

In the example shown in Figure 4-6, two DHCP servers are located on two different subnets. To configure the Accelar routing switch to forward the copies of the BootP/DHCP packets from the end station to both servers, specify the routing switch (10.10.1.254) as the agent address. Then enable DHCP to each of the DHCP servers by entering 10.10.2.1 and 10.10.3.1 as the server addresses.

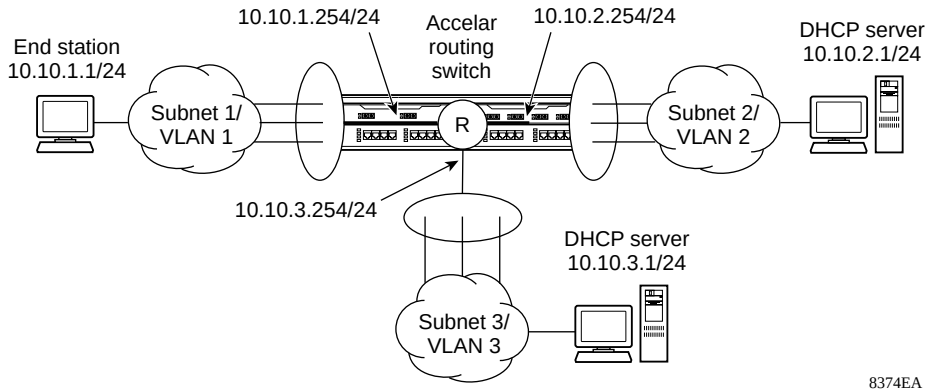

Figure 4-6. Configuring Multiple BootP/DHCP Servers

UDP Broadcast Forwarding

Some network applications, such as the NetBIOS name service, rely on a User Datagram Protocol (UDP) broadcast to request a service or locate a server for an application. If a host is on a network, subnet segment, or VLAN that does not include a server for the service, UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. Users work around this by forwarding the broadcasts to the server through physical or virtual router interfaces. Most routers allow configuring the interfaces to forward certain classes of broadcasts to the network subnet or directly to the server's IP address.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface out other router IP interfaces as a rebroadcast or to a configured IP address.

- If the address is that of a server, the packet will be sent as a unicast packet to this address.
- If the address is that of an interface on the router, the frame will be rebroadcast.

UDP Forwarding Operation

The basic steps for setting up UDP broadcast forwarding are:

1. Enter protocols into a table.
2. Create policies (protocol/server pairs).
3. Assemble these policies into lists or profiles.
4. Apply the list to the appropriate interfaces.

When a UDP broadcast is received on a router interface, in order to be considered forwarding, it must meet the following criteria:

- Be a MAC-level broadcast
- Be an IP limited broadcast
- Be for the specified UDP protocol
- Have a TTL value of at least 2

For each ingress interface and protocol, the policy specifies how the UDP broadcast is retransmitted: to a unicast host address or to a broadcast address.

Reverse Address Resolution Protocol (RARP)

Reverse Address Resolution Protocol (RARP) is a protocol used by some devices to obtain an IP address by providing their MAC layer address information to a RARP server. In previous versions of Accelar software, RARP was broadcast along with ARP and IP on all ports associated with an IP protocol-based or port-based VLAN. Therefore, it was not possible for a host to reach a RARP server outside the IP VLAN to get its IP address.

RARP has the format of an Address Resolution Protocol (ARP) frame but its own Ethernet type (8035). This makes it possible for RARP to be removed from the IP protocol-based VLAN definition and treated as a standalone protocol. By doing this, the concept of a RARP protocol-based VLAN is created.

A typical network topology provides desktop switches in wiring closets with one or more trunk ports extending to one or more data center switches where attached servers provide file, print, and other services. Using this new functionality, all ports in a network requiring the services of a RARP server could be defined as potential members of a RARP protocol-based VLAN. All tagged ports and data center RARP servers would be defined as static or permanent members of the RARP VLAN. Therefore, a desktop host would broadcast a RARP request to all other members of the RARP VLAN. In normal operation, these members would include only the requesting port, tagged ports, and data center RARP server ports. Because all other ports are potential members of this VLAN and RARP is only transmitted at bootup, all other port VLAN memberships would have expired. With this feature, one or more centrally located RARP servers could extend RARP services across traditional VLAN boundaries to reach desktops globally.

Virtual Router Redundancy Protocol (VRRP)

End stations are often configured with a static default gateway IP address. Loss of the default gateway router can have catastrophic results. Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure that can occur when the single static default gateway router for an end station is lost. It introduces the concept of a virtual IP address (transparent to users) shared between two or more routers connecting the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides a dynamic default gateway redundancy in the event of a failover.

The VRRP router controlling the IP address(es) associated with a virtual router is called the master router and forwards packets to these IP addresses. The election process provides a dynamic transition of forwarding responsibility if the master becomes unavailable.

In the configuration illustrated in Figure 4-7, the first three hosts install a default route to the virtual router 1 IP address and the other three hosts install a default route to the virtual router 2 IP address. This configuration not only has the effect of load sharing the outgoing traffic, it also provides full redundancy. If either router fails, the other router assumes responsibility for both addresses.

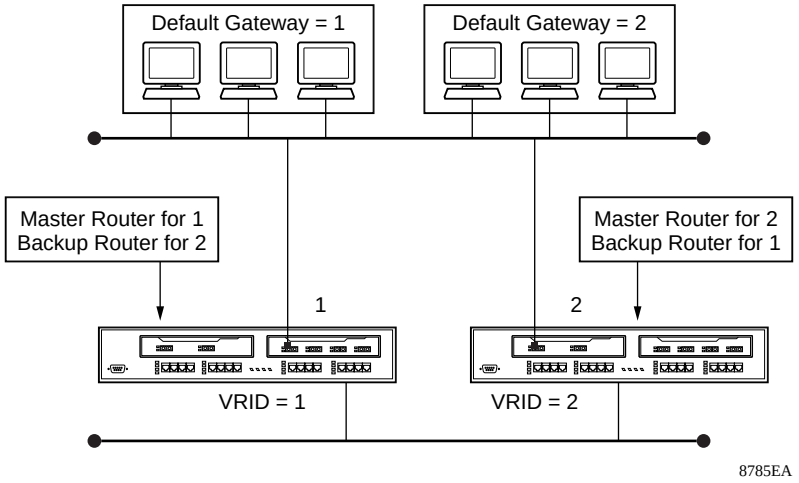

Figure 4-7. Example of VRRP in a Network

With version -A hardware, four VRRP interfaces (isolated routing ports *and* VLANs) are allowed per Accelar switch and all virtual router IDs (VRIDs) must be unique. Version -B hardware supports 256 VRRP interfaces per switch.

VRRP uses the following terms:

- VRRP router—a router running the VRRP protocol
- virtual router—an abstract object acting as the default router for one or more hosts, consisting of a virtual router ID and a set of addresses

- IP address owner—the VRRP router that has virtual router IP addresses as real interface addresses (This router is the one that will respond to packets sent to this IP address.)
- Primary address—an address selected from the real addresses and used as the source address of packets sent from the router interface
- Virtual router master—the router assuming responsibility for forwarding packets sent to the IP address associated with the virtual router and answering ARP requests for these IP addresses

When a VRRP router is initialized, if its priority is 255 (meaning that the router owns the associated VRRP addresses), it sends a VRRP advertisement, broadcasts an ARP request containing the virtual router MAC address for each IP address associated with the virtual router, and transitions to the master state. If the priority is not 255, the router transitions to the backup state.

In backup state, a VRRP router monitors the availability and state of the master router. It does not respond to ARP requests and must discard packets with a MAC address equal to the virtual router MAC address. It does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, it transitions back to the initialize state. If the master router goes down, the backup router sends the VRRP advertisement and ARP request described in the preceding paragraph and transitions to the master state.

In master state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. It responds to ARP requests for these IP addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and only accepts packets addressed to IP addresses associated with the virtual router if it is the IP address owner.

If a shutdown occurs, the VRRP router in master state sends a VRRP advertisement with priority of 0 and transitions to the initialize state. If an advertisement timer fires, the router sends an advertisement. If an advertisement is received with a 0 priority, the router sends an advertisement. If the priority is greater than the local priority or if it is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address, the router transitions to the backup state. Otherwise, it discards the advertisement.

Dynamic IP Routing Protocols

Unlike static IP routing, where a manual entry must be made in the routing table to specify a routing path, dynamic IP routing uses a “learning” approach to determine the paths and routes to other routers.

Accelar routing switches route dynamically in two ways:

- Routing Information Protocol (RIP) (this page)
- Open Shortest Path First (OSPF) Protocol (page 4-16)

Routing Information Protocol (RIP)

In a routed environment, routers communicate with one another to keep track of available routes. Routers can learn about available routes dynamically using the Routing Information Protocol (RIP). The Accelar routing software implements standard RIP for exchanging TCP/IP route information with other routers.

RIP uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. Each router “advertises” routing information by sending a routing information update every 30 seconds. If a router does not receive an update from another router within 90 seconds, it marks the routes served by the nonupdating router as being unusable. If no update is received within 240 seconds, the router removes all routing table entries for the nonupdating router.

Accelar switches also support RIPv2, which uses multicasting instead of broadcasting and supports variable length subnet masks (VLSM).

RIP is known as a distance vector protocol. The vector is the network number and next hop, and the distance is the cost associated with the network number. RIP identifies network reachability based on cost, and cost is defined as hop count. One hop is considered to be the distance from one router to the next. This cost or hop count is known as the *metric* (Figure 4-8).

A directly connected network has a metric of zero. An unreachable network has a metric of 16. Therefore, the highest metric between any two networks can be 15 hops or 15 routers.

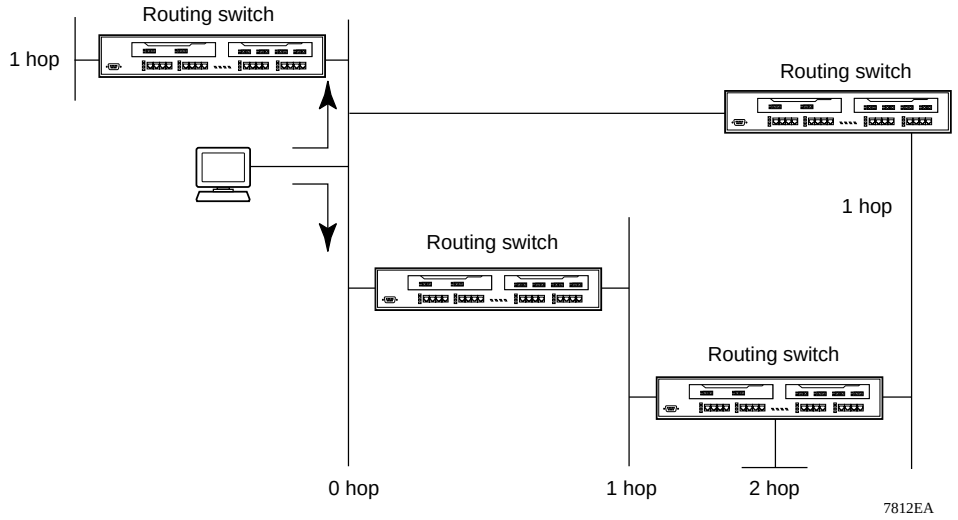

Figure 4-8. Hop Count or Metric in RIP

Open Shortest Path First (OSPF) Protocol

The Open Shortest Path First (OSPF) Protocol is an interior gateway protocol (IGP) intended for use in large networks. Using a link state algorithm, OSPF exchanges routing information between routers in an autonomous system. Routers synchronize their topological databases. Once the routers are synchronized and the routing tables are built, the routers will flood topology information only in response to some topological change. For OSPF, the “best” path to a destination is the path that offers the least cost metric delay. In OSPF, cost metrics are configurable, allowing you to specify preferred paths.

Routers keep a table of “reachability information” containing a list of networks and routers. The table is maintained with current information via a regular exchange of hello packets.

OSPF is a *link-state protocol.* A router running a link-state protocol periodically tests the status of the physical connection to each of its neighbor routers and sends this information to its other neighbors. A link-state protocol does not require each router to send its entire routing table to its neighbors. Instead, each OSPF router floods only link-state change information throughout the autonomous system (or area, if the AS is divided into areas). This process is referred to as the synchronization of the routers' topological databases.

With the link information, each router builds a *shortest-path tree* with itself as the root of the tree. It then can identify the shortest path from itself to each destination and build its routing table.

OSPF Addresses and Variable-Length Masks

A destination in an OSPF route advertisement is expressed as an IP address and a variable-length mask. Taken together, the address and the mask indicate the range of destinations to which the advertisement applies.

The ability to specify a range of networks allows OSPF to send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 128.185.0.0 with a mask of 0xffff0000 describes a single route to destinations 128.185.0.0 to 128.185.255.255.

OSPF Neighbors

OSPF *neighbors* are any two routers that have an interface to the same network. In each OSPF network, routers use the Hello Protocol to discover their neighbors and maintain neighbor relationships. On a broadcast or point-to-point network, the Hello Protocol dynamically discovers neighbors; however, on a nonbroadcast multiaccess network, you must manually configure neighbors.

The Hello Protocol ensures that communication between neighbors is bidirectional. Periodically, OSPF routers send out hello packets over all interfaces. Included in these hello packets is the following information:

- The router's priority
- The router's Hello Timer and Dead Timer values
- A list of routers that have sent this router hello packets on this interface
- The router's choice for designated router and backup designated router

Bidirectional communication is determined when one router sees itself listed in the neighbor's hello packet.

Neighbor Adjacencies

Neighbors may form an *adjacency* for the purpose of exchanging routing information. When two routers form an adjacency, they go through a process called *database exchange* to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent. From this point on, only routing change information is passed between the adjacencies, thus conserving bandwidth.

All routers connected by a point-to-point network or a virtual link will always form an adjacency. Also, every router on a multiaccess network forms an adjacency relationship with the designated router and the backup designated router.

Designated Routers

To further reduce the amount of routing traffic, the Hello Protocol elects a *designated router* and a *backup designated router* on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information with each other (which on a large network can mean a lot of routing protocol traffic), all routers on the network form adjacencies with the designated router and the backup designated router *only* and send link state information to them. The designated router then redistributes the information from each router to every other router.

The Hello Protocol always elects a backup designated router along with the designated router. This router takes over all of the designated router's functions if that router fails.

OSPF Areas

OSPF routers reduce and restrict the amount of internal and external routing information that is flooded through the Autonomous System (AS) by dividing the AS into *areas*. Each area has a unique ID number. (ID 0.0.0.0 is always reserved for the OSPF backbone.)

Two specialized types of areas are *stub areas* and *not so stubby areas* (NSSAs), both designed to preserve router resources. A stub area does not receive advertisements for external routes, which reduces the size of the link state database. A stub area has only one area border router; packets destined outside the area are routed to that area border exit point, examined by the area border router, and forwarded to a destination.

An NSSA also prevents the flooding of AS-External-Link State advertisements into the area by replacing them with a default route. The added feature of NSSAs is the ability to import small stub (non-OSPF) routing domains into OSPF. Like stub areas, NSSAs are at the edge of an OSPF routing domain. The non-OSPF routing domains are attached to the NSSAs, forming the NSSA transit areas. Providing the addressing scheme of non-OSPF stub domains permits the NSSA border to also perform manual aggregation.

OSPF Router Types

Routers deployed in an OSPF network can take on different roles depending on how they are configured. Table 4-1 provides a brief description of each possible router role.

Table 4-1. Router Classifications

Router Type	Description
AS Boundary Router (ASBR)	A router attached at the edge of an OSPF network is considered an AS Boundary Router (ASBR). An ASBR generally has one or more interfaces that run an Inter-Domain Routing Protocol (IDRP) such as BGP. In addition, any router distributing static routes or RIP routes into OSPF is considered an ASBR. The ASBR forwards routes learned from IDRP into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area Border Router (ABR)	A router attached to two or more areas inside an OSPF network is considered an Area Border Router (ABR). ABRs play an important role in OSPF networks by limiting the amount of OSPF information that gets disseminated.
Internal Router (IR)	A router that only has interfaces within a single area inside an OSPF network is considered an Internal Router (IR). Unlike ABRs, IRs have topological information only about the area in which they are contained.

Table 4-1. Router Classifications (continued)

Router Type	Description
Designated Router (DR)	In a broadcast network, such as an Ethernet network that has more than one router locally attached, a single router is elected to be the Designated Router (DR) for that broadcast network. A DR assumes the responsibility of making sure all routers on the broadcast network are in synchronization with one another. The DR is also responsible for advertising that network to the rest of the autonomous system.
Backup Designated Router (BDR)	In a broadcast network, such as an Ethernet network, a Backup Designated Router (BDR) is elected in addition to the Designated Router (DR). The BDR can assume essentially the same responsibilities as the DR; if the DR fails, the BDR assumes the role of the DR in the broadcast network.
Other Router (OR)	In a broadcast network, such as an Ethernet network, any router not elected to be a Designated Router (DR) or Backup Designated Router (BDR) is considered to be an Other Router (OR).

AS External Routes

OSPF considers the following routes to be *AS external (ASE)* routes:

- A route to a destination outside the AS
- A static route
- A default route
- A route derived by RIP
- A directly connected network not running OSPF

IP Policies

IP policies are designed to optimize the use of a routing table and allow better control over an otherwise-dynamic routing protocol.

IP *accept policies* govern the addition of new RIP- or OSPF-derived routes to the routing tables. When RIP or OSPF receives a new routing update, it consults its accept policies to validate the information before entering the update into the routing tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).

IP *announce policies* govern the propagation of RIP or OSPF routing information. When preparing a routing advertisement, RIP or OSPF consults its announce policies to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure).

The flow of routing information between the network, the protocols, and the routing table manager is controlled by *routing information policies*. Each time a routing update arrives from a remote router, the following steps occur:

1. The protocol receiving the route consults an *accept policy* to determine whether to forward the route or drop the route.
2. The protocol consults an *announce policy* to determine whether or not to advertise the route to the network.

Policies in Accelar Switches

Accept and announce policies are configured for the Accelar routing switch based on the selected protocol (OSPF or RIP).

A policy is made up of three parts: matching criteria, set parameters, and action. The matching criteria are used to decide whether or not a policy should be applied to a certain route. Once a policy is selected for a route, the set parameters are used to construct the route advertisement only if the action is announce.

Announce policies enable a user to selectively announce routes. Announce policies alter the routing information learned by the routers in a particular routing domain. OSPF announce policies are applied for non-OSPF routes in an Autonomous System Boundary Router (ASBR). Only an ASBR advertises the external route information into the OSPF domain. If no policies are configured or no matching policy exists for a given route, the default behavior is applied; that is, OSPF ignores the external route information.

OSPF accept policies are applied whenever the OSPF engine computes the external routes due to a topology change or an external link-state advertisement (LSA). If there are no policies configured or no matching policy is found for a given route, the default behavior is applied; that is, the external route is included in the routing table.

The method in which OSPF applies accept and announce policies to routing information differs somewhat from the way RIP handles policies. OSPF link-state LSAs are received and placed in the link state database (LSDB) of the router. The information in the LSDB is also propagated to other routers in the OSPF routing domain. According to the OSPF standard, all routers in a given area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them to other routers.

To accomplish this, OSPF accept and announce policies act in the following manner:

- OSPF accept policies control which OSPF non-self-originated external routing information is processed. The accept policies control only what the local router uses; they do not affect the propagation of OSPF internal and OSPF non-self-originated external information to other routers.
- OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn but only with regard to the local router's self-originated information.

RIP announce policies are applied while sending a RIP update. The policy information is used to announce the route to other routers in the RIP routing domain. If no policies are configured or no matching policy exists for a given route, the default behavior is applied; that is, RIP-learned routes will be announced and all non-RIP routes will be ignored.

RIP accept policies are applied whenever the router receives a RIP update. The policy is used to selectively accept routes from the RIP update. If no policies are configured or no matching policy exists for a given route, the default behavior is applied; that is, the route is included in the routing table.

For specific information about creating policies in Accelar switches, refer to *Reference for Accelar Management Software Routing Operations* when using Device Manager or *Reference for the Accelar 1000 Series Command Line Interface* when using the CLI.

Chapter 5
IP Multicasting

This chapter describes the IP Multicast features supported on an Accelar 1000 Series routing switch. IP multicasting provides services such as the delivery of information to multiple destinations with a single transmission and the solicitation of servers by clients. These services benefit applications such as audio/video conferencing, dissemination of data to multiple recipients, network meetings, and so on.

Two protocols are used to enable IP multicasting in Accelar switches:

- Internet Group Management Protocol (IGMP) is used by hosts to report their multicast group memberships to neighbor multicast routers.
- Distance Vector Multicast Routing Protocol (DVMRP) is used between routers to exchange multicast routing information.

Multicast Host Groups

IP Multicast is a method for addressing, routing, and delivering a datagram to a collection of receivers called a *host group*.

Host groups can be permanent or transient:

- A *permanent host group* has a well-known, administratively assigned IP Multicast group address. The address, not the membership, is permanent and defines the group. A permanent host group can consist of zero or more members.
- A *transient host group* exists only as long as it has members that need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

Any host system on any IP network can send a message to a multicast group using the group's IP Multicast address. To receive a message addressed to a multicast group, however, the host must be a member of the group and must reside on a network where that group is registered with a local multicast router.

An IP multicasting host group can consist of zero or more members and places no restrictions on its membership. Host members can reside anywhere; they can join and leave the group at any time; and they can be members of more than one group at the same time. In order to receive a multicast message from a host group, a host must be a member of the group. However, anyone can send a multicast datagram: A host does not need to be a member of a group to send a multicast message to its members.

In general, hosts that are members of the same group reside on different networks. However, a range of multicast addresses (224.0.0.*x*) is reserved for groups that are locally scoped. All message traffic for these hosts typically remains on the local network. Hosts that belong to a group in this address range and that reside in different networks will not receive each other's message traffic.

Note: A special set of filters apply to multicast packets. The user is capable of creating deny or accept filters to configure which sources can receive and send data.

Multicast Addresses

Each host group is assigned a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP Version 4 multicast address is a Class D address (the high-order bits are set to 1110) from 224.0.0.0 to 239.255.255.255.

The block of addresses from 224.0.0.1 to 224.0.0.255 is reserved for routing protocols and other low-level protocols. Multicast routers will not forward datagrams with addresses in this range because the TTL is usually set to 1.

Internet Group Management Protocol

The Internet Group Management Protocol (IGMP) has the following characteristics:

- IGMP allows a host to register group memberships with the local querier router to receive any datagrams sent to this router and targeted to a group with a specific IP Multicast address.
- IGMP allows a router to learn the existence of group members on its directly attached networks. The router periodically sends a general query message to each of its local networks. Any host that is a member of any multicasting group identifies itself by sending a response.

IGMP Concepts and Terminology

IGMP is a protocol used by IP Multicast routers to learn the existence of host group members on their directly attached subnets. It allows hosts to communicate their desired group memberships to their local querier router to receive any datagrams sent to this router and targeted to a group with a specific IP Multicast address. A router communicates with the hosts on a local network by sending IGMP queries. Hosts respond by issuing IGMP reports.

This section covers the following topics:

- IGMP queries (this page)
- IGMP host reports (page 5-4)
- Host leave messages (page 5-5)

IGMP Queries

A router running IGMP periodically sends host membership queries (also known as General Queries) to its attached local networks. The Accelar switch supports both version 1 queries and version 2 queries.

If multiple IGMP routers exist on the network, one router is designated to send queries, using the following rules:

- Choose a router that generates version 1 queries over a router that generates version 2 queries.
- Choose the router with the lowest IP address when running version 2. In version 1, the router with the highest IP address becomes the IGMP querier.

IGMP Host Reports

A host that receives a membership query from a local router can respond with a *host membership report*, one report for each joined multicast group. A host that receives a query delays its reply by a random interval and listens for a reply from any other host in the same host group. Consider a network that includes two host members—host A and host B—of the same multicast group. The router sends out a host membership query on the local network. Both host A and host B receive the query and listen on the network for a host membership report. Host B's delay time expires first, so it responds to the query with a membership report. Hearing the response, host A does not send a report of its own for the same group.

Each query from a router to a host includes a Maximum Response Time field. IGMP inserts a value—*n* —into this field specifying the maximum time in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and *n* tenths of a second for the period that it waits before sending a response. This is true for IGMP Version 2 only. For IGMP Version 1, this field is set to 0 but defaults to a value of 100, that is, 10 seconds.

If at least one host on the local network specifies that it is a member of a given group, the router will forward to that network all datagrams bearing the group's multicast address.

Upon initialization, the host may immediately issue a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same way it accepts requested reports.

Once in a steady state, hosts and routers communicate in a way that minimizes the exchange of queries and reports. The designated routers set up a path between the IP Multicast stream source and the end stations and periodically query the end stations about whether or not to continue participation. As long as any client continues to participate, all clients, including nonparticipating end stations on the switch port, receive the IP Multicast stream.

Host Leave Messages

When a version 2 host leaves a group and it was the host that issued the most recent report, it issues a *host leave message*. The multicast router on the network issues a group-specific query to determine whether there are other group members on the network. If no host responds to the query, the router assumes that no members belonging to that group exist on that interface.

Accelar IGMP Implementation

The Accelar switch supports IGMPv1 and IGMPv2 registration protocols and will generate IGMP queries on all subnets and interfaces for which IP multicasting is enabled. Multicast frames that arrive from an interface are forwarded on all interfaces/subnets on which IGMP reports have been received for the multicast group indicated in the destination IP address.

Multicast packets forwarded within the same VLAN remain unchanged and packets are not forwarded to networks with no members of the multicast group indicated in the destination IP address.

Multicast routing can be enabled and disabled on an interface basis. If multicast routing is disabled on an interface, IGMP queries are not generated. If the switch is in IGMP router behavior mode, IGMP snooping is not configurable.

IGMP Snooping

When used as a switch, Accelar 1000 series products support IGMPv1 and IGMPv2 to prune group membership per port within a VLAN. This feature is called IGMP snooping.

The IGMP snooping feature allows the user to optimize the multicast data flow for a group within a VLAN only to those ports that are members of the group. The switch listens to group reports from each port and builds a database of multicast group members per port. It suppresses the reports heard by not forwarding them out to other hosts, forcing the members to continuously send their own reports. The switch relays group membership from the hosts to the multicast routers. It forwards queries from multicast routers to all port members of the VLAN. Furthermore, it multicasts data only to the participating group members and to the multicast routers within the VLAN.

The number of multicast groups per VLAN limits the number of available VLANs on the switch for ARU2/QUID4. In QUID5, multiple groups share hardware resources required for VLANs.

IGMP Proxy

If the Accelar switch receives multiple reports for the same multicast group, it does not transmit each report to the multicast upstream router. Instead, it forwards only the first report. If there is new information that another multicast group has been added or that a query has been received since the last report was transmitted upstream, then the report will be forwarded onto the multicast router ports. This feature is known as IGMP proxy.

Distance Vector Multicast Routing Protocol

DVMRP is a distance vector type of multicast routing protocol. It advertises shortest-path routes to multicasting *source networks*—that is, any network containing hosts that have the capability to issue multicast datagrams. (In this respect, DVMRP is the opposite of RIP, which advertises routes to destination networks.) Coupled with IGMP, membership for a multicast stream is learned from both the routers and directly attached hosts.

DVRMP constructs a different distribution tree for each source and its destination host group. Each tree is a minimum multicast tree from the multicast source at the root of the tree to all multicast receivers as leaves of the tree. The distribution tree provides a shortest path between the source and each multicast receiver in the group, based on the number of hops in the path. A tree is constructed on demand, using a broadcast and prune technique, when a source begins to transmit messages to a multicast group.

DVRMP assumes that initially every host on the network is part of the multicast group. The designated router on the source subnet (the router that has been selected to handle routing for all host on the subnet) begins transmitting a multicast message to all adjacent routers. Each of these routers then selectively forwards the message to downstream routers until the message is eventually passed to all multicast group members.

In the selective forwarding process during the formation of the multicast tree, when a router receives a multicast message, it checks its unicast routing tables to determine the interface that provides the shortest path back to the source. If that was the interface over which the multicast message arrived, the router enters some state information to identify the multicast group in its internal tables and forwards the multicast message to all adjacent routers except the one that sent the message. If the interface was not the one receiving the multicast message, the message is discarded. This mechanism, called reverse path forwarding, ensures that there will be no loops in the tree and that the tree will include the shortest path from the source to all recipients.

DVRMP actually uses a forwarding process that is even more selective by relying on specific information that is provided by the unicast routing protocol. This enhancement potentially results in a reduction in the number of flooding messages required to construct the distribution tree.

The pruning feature of the protocol eliminates branches of the tree that do not lead to any multicast group members. The IGMP running between hosts and their immediately neighboring multicast routers is used to maintain group membership data in the routers. When a router determines that no hosts beyond it belong to the multicast group, it sends a prune message to its upstream router. Routers update source and destination group state information in their tables to reflect which branches are eliminated from the tree, resulting in a minimum multicast tree. If a router later learns of new group memberships from the hosts or downstream routers, it sends a graft message upstream to retract the prune sent earlier.

After the multicast tree is constructed, it is used to transmit multicast messages from the source to multicast members. Each router in the path forwards messages over only those interfaces that lead to group members. Because new members can join the group at any time and these members may depend on one of the pruned branches to receive the transmission, DVRMP periodically reinitiates the construction of the multicast tree.

DVMRP Concepts and Terminology

DVMRP is a multicasting protocol that provides a mechanism for routers to propagate multicast datagrams in a manner that minimizes the number of excess copies sent to any particular network.

This section covers the following topics:

- Neighbor connections (page 5-8)

- Source route advertisements (page 5-8)
- How DVMRP chooses a route (page 5-9)
- Routing table (page 5-9)
- Shortest-path trees (page 5-10)

Neighbor Connections

In a DVMRP environment, *neighbors* are multicasting routers that have an interface to the same network.

At startup, a DVMRP multicasting router performs the following tasks:

- Initializes its routing table with information on all of its local networks.
- Learns the existence of its neighbors by sending a probe for all routes on each of its multicast interfaces.
- Receives reports from its neighbors containing the routing information (including route costs).

Source Route Advertisements

A *source network* is any network containing hosts that have the capability to issue multicast datagrams. DVMRP advertises shortest-path routes to multicasting source networks. In this respect, DVMRP is the opposite of RIP, which advertises routes to destination networks.

Periodically, each multicasting router issues full or partial routing information on each DVMRP circuit, using DVMRP report messages. This routing information represents the sending router's cost to reach the specified source network. (The cost is the sum of the hop metrics along the shortest path to the given source network.)

Upon receiving a DVMRP report from another router, DVMRP reexamines its routing table to determine if the shortest path information needs updating. Specifically, DVMRP looks in the routing table for an entry describing a route to the same source network. If one exists, DVMRP compares the cost of the two routes and stores the route with the lower cost in its routing table.

A router will not send route reports on an interface until it knows (by means of received probes or reports) that it has a neighboring multicast router on that interface. It will continue to send probes periodically on an interface.

How DVMRP Chooses a Route

Each DVMRP interface is configured with a metric that indicates the cost of the hop. A router that receives multiple route reports for the same multicasting source network performs the following tasks:

- Compares the cost specified in each (based on the metric field).
- Stores information from the report with the lowest cost in its routing table.

A route metric is the sum of all the interface (hop) metrics from a given route source to a given router. After a next-hop neighbor has been declared for a route, the route updates received from that neighbor for that route take precedence until either the route times out or another router advertises a better metric for that route.

Routing Table

Table 5-1 shows the principal items in a routing table entry.

Table 5-1. Parts of a Routing Table Entry

Item	Description
Source subnet address and mask	The network address and mask that identify the source for which this entry contains multicast routing information.
Upstream neighbor	The address of the upstream neighbor from which IP datagrams from these sources are received.
Interface	The value of the interface index on which IP datagrams sent by these sources are received.
Metric	The distance in hops to the source subnet.
Expiration Time	The maximum amount of time (in timeticks) remaining before this entry will be aged out.

Note that the source subnet and the previous-hop router in the DVMRP routing table are the opposite of the destination subnet and next-hop router in a RIP routing table.

Using this information, the router performs the following tasks:

- Receives a multicast datagram and determines if the datagram has arrived on the interface that is on the shortest path to the source network.
- Drops the datagram if it has not arrived on the shortest-path interface.
- Floods the multicast datagram to all active downstream DVMRP neighbors.

Shortest-Path Trees

Route information used by DVMRP is independent of any other routing information used by the router (for example, routes provided by RIP). The purpose of this routing information is to create a shortest-path tree entry in the routing table for the propagation of multicast datagrams.

The shortest-path tree entry indicates the interface that provides the shortest path to the network that is the source of the multicast datagram. A shortest-path tree also indicates those interfaces that are on the shortest path to that source network from a neighboring router.

In IGMP version 2, neighboring routers have the same metric to a given source network.The router with the lower IP address is responsible for propagating multicast traffic originating from that source network onto the network or tunnel that is common to these neighboring routers.

A network is considered a leaf network if it has no dependent downstream neighbors for a source.

Accelar Implementation of DVRMP

DVRMP in Accelar switches fully supports multiaccess networks. The forwarding entries for the receivers on multiaccess networks are port based rather than network based. Therefore, on a multiaccess network, data will not be received by any ports other than the ones interested in the data. In other words, IP Multicast routing is supported on ports with port-based or IP subnet-based VLANs enabled.

The DVRMP router will listen to all IGMP host membership reports, even if it is not the designated querier and will keep a local group database of every host membership reporter.

When a multicast stream (UDP packets) first enters the switch, if DVRMP is enabled for the interface, then DVRMP will process this packet as necessary and create a hardware cache entry to handle subsequent UPD packets for the same multicast destination. The packets are discarded if there are no members; otherwise, they are forwarded.

The Accelar implementation does not support DVMRP tunneling.

Chapter 6
IP Filtering

This chapter describes using IP filtering as supported on an Accelar 1000 Series routing switch. These filters are used to manage traffic and, in some cases, to provide security. Each filter set includes match conditions (a set of filter records defining the match criteria) and actions (forward, drop, prioritize, mirror) to be performed when a match condition is satisfied. IP filtering support was added with the version -A hardware and enhanced with version -B hardware.

IP filters apply to all routed IP packets to be forwarded through the routing switch on specified ingress ports. The filters are applied to the switch ingress ports with a default action to forward or drop. All packets not matching any filter are forwarded or dropped, depending on the port's default action. In version -A hardware, the port default action was to forward; in version -B hardware, the port default action can be either forward or drop.

Two types of filters can be applied: traffic filters (source or destination) or global filters. Filters are applied to a port using filter sets, and actions are assigned when applying a filter set to a port. The actions of individual filters can overwrite the default actions of the port.

Filter Characteristics

Filters on Accelar routing switches have the following characteristics and requirements:

- Up to 1024 filter IDs can be defined, including source/destination and global filters. These filter IDs are system generated.
- Up to 128 filter sets can be defined, including source/destination and global filter sets.

- A collection of source/destination filters is defined in a set, and the set is applied to a port or group of ports. Multiple sets can be assigned to any given port.
- A collection of global filters is defined in a global set (not exceeding eight per set), and the set is applied to a port or group of ports. Multiple sets may be applied to a given port or set of ports, but the maximum number of global filters that can be enabled on a given port set is eight.

Beginning with this release, filter counters are maintained for all active filters. Each time an active filter is "hit" by a packet, its counter is incremented by one. These counters are maintained chassis-wide and may be viewed or reset administratively at any time. The counters may also be reported to a syslog host.

Source and Destination Filters

Source and destination filters (traffic filters) instruct a router interface to selectively handle specified IP traffic. You determine which packets receive special handling based on information in the packet headers. Using traffic filters, you can reduce network congestion and control access to network resources by blocking, forwarding, or prioritizing specified traffic on an interface. You can apply multiple traffic filters to a single interface.

Source filters must specify a source IP address and mask, and they may optionally specify a destination IP address and mask. Destination filters must specify a destination IP address and mask, and they may optionally specify a source IP address and mask. The minimum mask length is eight bits. A source or destination filter can cause the following actions to be applied to a packet that matches the filter record:

- Forward the packet when the filter is applied with a forward action.
- Drop the packet when the filter is applied with a drop action.
- Mirror the packet to the defined mirror port.
- Elevate priority to high.
- Elevate the IEEE VLAN tag priority (0 through 7). If this priority bit field is greater than the priority field currently in the IEEE tag in the packet header, this priority field will be loaded into the packet header.
- Manipulate the type of service (TOS) bit.

Beginning with this release, configuring source or destination filters for non local routes will, as needed, cause corresponding routing entries to be created and maintained in the routing table so that the effectiveness of these filters is maintained.

Global Filters

Global filters may specify a source IP address and mask, a destination IP address and mask, both of these, or neither of these. Global filters have the following characteristics:

- No minimum or maximum mask length exists.
- Up to eight global filters can be applied on any given set of ARU ports. A set includes four 10/100 mb/s ports or one Gigabit port, each of which can accommodate eight global filters.
- A global filter can cause the same actions described above for source/ destination filters.

Filter Configuration

Matching criteria for filters in Accelar switches can be any of the following:

- Destination address or address range
- Source address or address range
- Exact IP protocol match (TCP, UDP, or ICMP)
- TCP or UDP port numbers
- TCP connections established from within the network only or bidirectional establishment allowed

Configurable actions are:

- Drop
- Forward
- Mirror
- High priority (within the switch fabric)
- TCP connect (prevent incoming TCP sessions)
- IEEE VLAN priority (increase 802.1Q tag priority)

Actions

Each filter has an action mode associated with it, which determines whether packets matching this filter are forwarded (routed) through the switch.

Each physical port on the Accelar switch has a default action of forward or drop associated with it. When the filtering action mode matches the port default action, the port default action is taken.

When the port default action is drop, a packet will be forwarded only if a matching filter was set with an action mode of forward. If a single match occurs with an action mode of forward, it does not matter how many matching filters are found with an action mode of drop; the frame will still be forwarded. That is, if a packet matches multiple filters and any one of them has an action mode of forward, the packet will be forwarded.

When the port mode is set to forward, a packet will be dropped only if a matching filter is found with a drop action mode. Again if a single match occurs with a drop action, it does not matter how many matching filters have forwarding actions; the packet will be dropped. If a packet matches multiple filters and any one of them has an action mode of drop, the packet will be dropped.

Table 6-1 indicates the forward/drop behavior of a port if multiple filter matches are found for a packet.

Table 6-1. Port Actions for Combinations of Matching Filters

Port Mode	Forward Filter Match	Drop Filter Match	Packet Action
Drop	Yes	No	Forward
Drop	Yes	Yes	Forward
Drop	No	No	Drop
Drop	No	Yes	Drop
Forward	Yes	No	Forward
Forward	Yes	Yes	Drop
Forward	No	Yes	Drop
Forward	No	No	Forward

Chapter 7
IPX Routing

The Accelar implementation of IPX supports four Ethernet frame formats:

- Ethernet II
- 802.2-LLC
- 802.3-RAW
- 802.3-SNAP

Frame translations from one frame format to another frame format are supported. However, the IPX host format must match the defined protocol frame format of the interface to which the host is connected.

In Accelar switches, IPX is supported on routed VLANs. IPX routing is not available on isolated routing ports.

IPX Protocol

The Internetwork Packet Exchange (IPX) Protocol is the Novell, Inc. adaptation of the Xerox Network System (XNS) Protocol.

IPX has the following characteristics:

- It is a *connectionless datagram* delivery protocol. *Connectionless* means that it does not need a channel established for packet delivery. A *datagram* is a unit of data that contains all the addressing information required for it to be delivered to its destination. The path or route that one datagram follows to reach a destination is independent of the path or route that another datagram may follow to reach the same destination.

- It does not guarantee the delivery of packets. Upper-layer protocols assume the responsibility for reliability. The upper-layer protocols that IPX uses are SPX and NCP.
- It uses the Internet Data Packet (IDP) format.

IPX is the network layer routing protocol used in the NetWare environment. The primary tasks of IPX are addressing, routing, and switching information packets from one location to another on a network. The network interface card (NIC) in a client provides network node addressing. IPX defines the *internetwork* and *intranode* addressing as follows:

- *Network addresses* form the basis of the IPX internetwork addressing scheme for sending packets between network segments. Every network segment of an internetwork is assigned a unique network address that routers use to forward packets to their final destinations. A network address in the NetWare environment consists of 8 hexadecimal characters. In the following example, *0x* indicates that this is a hexadecimal number, and *n* is any hexadecimal character:

  0xnnnnnnnn

- *Socket numbers* are the basis for an IPX intranode address, that is, the address of an individual entity within a node. They allow a process (for example, RIP or SAP) to distinguish itself to IPX. To communicate on the network, the process must request a socket number. Any packets that IPX receives addressed to that socket are then passed on to the process within the node.

The most basic network configuration consists of a single client and a file server. Major server characteristics are:

- Server name — a logical entity used to aid in managing services and devices
- Internal network address
- MAC address — a 6- byte PROM based in the NIC
- Server type — automatically assigned by Novell based on services offered (such as print server, file server, or application server)

An internetwork is created when two or more networks are joined together by a router. In an internetwork, each network segment maintains all the properties of a multiple server network with the connection device (the router) distinguishing one segment from another by its network address.

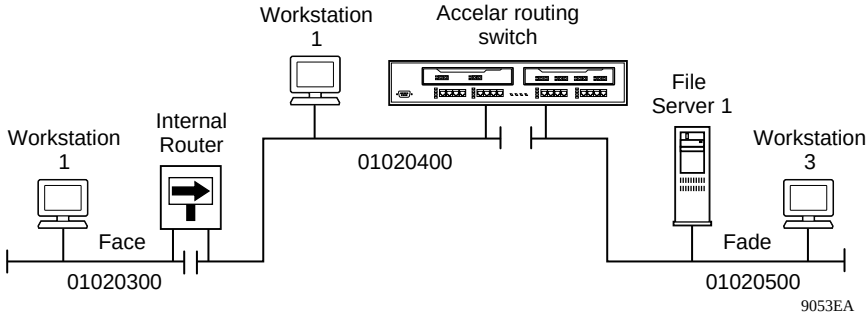

Figure 7-1. IPX Internetwork

In Figure 7-1, networks 01020300, 01020400, and 01020500 are joined by routers. Each segment has its own unique network address, which is assigned for each interface of the routing device and used to identify each network segment. Routing devices can be file servers acting as a router (internal routers) or dedicated routers, such as an Accelar routing switch.

Network layer addresses use hexadecimal characters and are one to eight digits in length. Valid addresses are 00000001 to FFFFFFFD; address FFFFFFFE is reserved for the default route. Network addresses, commonly called external addresses, are assigned to NICs during installation. In addition, most Novell routing devices are also assigned an internal address, the logical network where all file server services reside.

Figure 7-1 shows a 3-segment Novell internetwork. Network 01020300 contains workstation 1 and the left half of the file server. Network 01020400 includes the right half of the file server, workstation 2, and the left side of the router. Network 01020500 consists of the right side of the router, file server 1 and workstation 3. The internal network addresses, FACE and FADE, are assigned to the file servers. Clients are not assigned network layer addresses but rather acquire them from the nearest routing device.

File servers provide shared services to network users. File servers enable network users to share common network services and devices as well as network applications loaded on the file server. All services are advertised by SAP using Split Horizon techniques. All services are tied to a network address, and you must

be able to reach the network address in order to access the services. Applications are statically stored on the server; when network users log on to the server and request to use an application, the application is downloaded and executed within the memory of the workstation.

Communication between a server and client require the use of numerous protocols. As data makes its way through the protocol stack of one device on the way to another device, the data is continuously enveloped by various protocol headers until it is sent to the destination device. When the destination device is reached, the headers that were added by the source device are sequentially stripped away.

IPX Packets

An IPX packet consists of two major parts: a 30-byte header and the data portion. The destination and source network, node, and socket addresses make up the IPX packet header. The header contains fields with the following information:

- Checksum
- Packet length (Total length can range from 30 to 64,535 bytes.)
- Transport control— indicating how many routers an IPX packet has passed through on its way to the destination network (When the value reaches 16, the packet is discarded.)
- Packet type —a 1-byte field (Novell-specified) to indicate the type of service such as RIP (1), SAP (4), NetBIOS (20), NCP (17), unknown (0)
- Destination network, destination node, and destination socket
- Source network, source node, and source socket
- Data

Figure 7-2 shows the contents of the header, and Figure 7-3 shows the contents of the four supported types of IPX packets.

Checksum = FFFF	2 bytes
Packet Length	2 bytes
Transport Control	1 byte
Packet Type	1 byte
Destination Network	4 bytes
Destination Node	6 bytes
Destination Socket	2 bytes
Source Network	4 bytes
Source Node	6 bytes
Source Socket	2 bytes
Data	

Figure 7-2. IPX Header

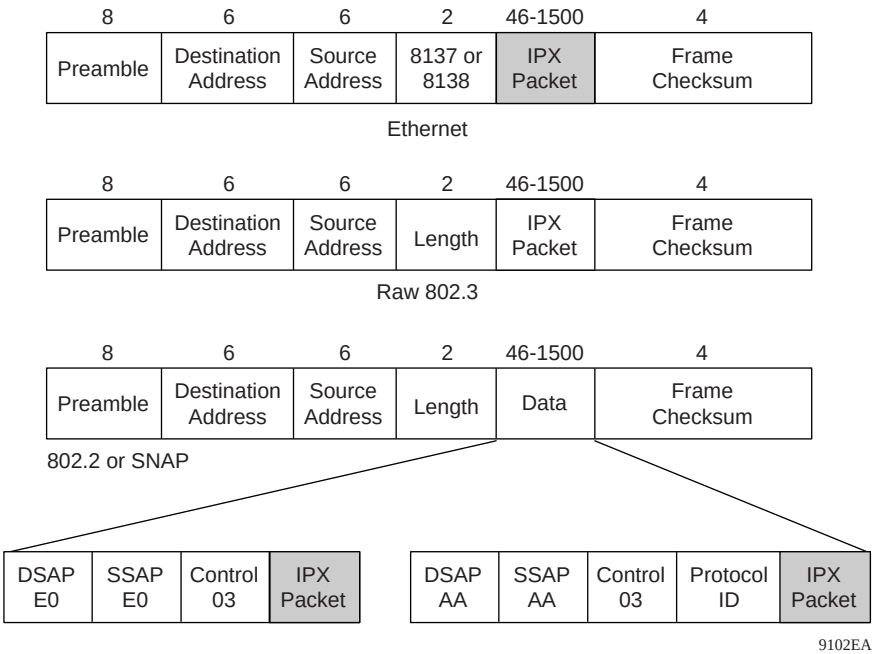

Figure 7-3. IPX Packets

IPX Traffic

IPX devices communicate through an IPX router. Figure 7-4 shows a typical example.

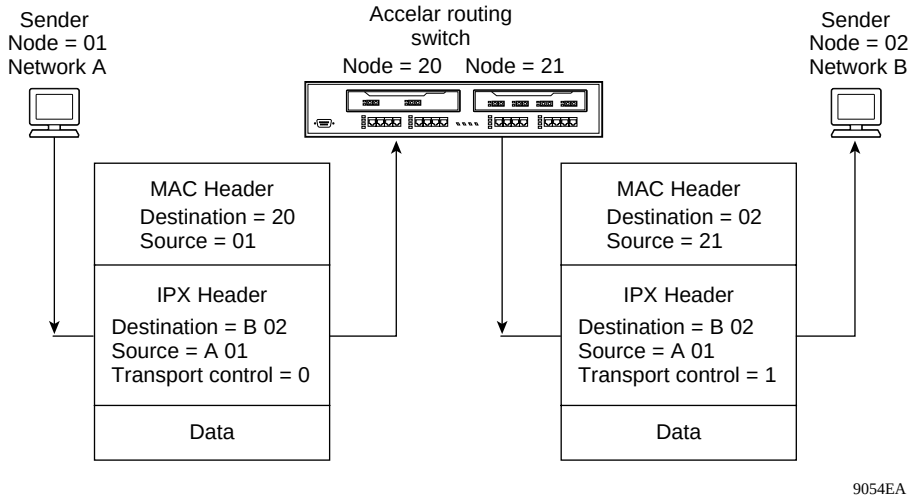

Figure 7-4. IPX Addressing

The sender and receiver can be any type of Novell device (for example: two workstations, two routers, two file servers, or a workstation and a server).The sender establishes a connection with a router on its own segment and then prepares its packet to send. The sending node places the destination node internetwork address (network address, node address, and socket number) in the destination address fields of the IPX header and its own internetwork address in the source address fields.

In the MAC header, it places the node address of the router that responded to its RIP request in the destination field, and its own node address in the source address field. It then sends the packet through the network media.

When the router receives the packet, if it is directly connected to the packet's destination segment, it places the destination node address from the IPX header in the destination address fields of the MAC header and increments the hop count in the transport control field by one. It then places its own node address into the source address field of the MAC header and transmits the packet.

If the router is not directly connected to the destination network segment, it passes the packet to the next router in the path to the destination node, by placing the next router node address into the destination address field of the MAC header. It places its own node address in the source address field, increments the hop count by one, and sends the packet to the next router.

When the packet arrives at the final destination, the MAC header is stripped out by the packet driver in the receiving node, the IPX header is stripped out by the IPX protocol, and the data is passed to the receiving node.

Accelar IPX Network-Layer Support

An Accelar router running IPX provides the following network-layer support:

- Dynamic routing of IPX packets
- Up to four IPX network addresses to an interface
- Routing Information Protocol (RIP)
- Service Advertisement Protocol (SAP)
- Static route support
- Default route support

Routing Information Protocol (RIP)

In a Novell internetwork, routers use RIP to exchange routing information. One RIP broadcast packet can contain up to 50 sets of network number information (routes). RIP packets do the following:

- Enable workstations to locate the fastest route to a network by broadcasting a route request
- Allow routers to update internal routing tables by broadcasting a route request
- Allow routers to respond to route requests from workstations and other routers
- Allow routers to perform periodic broadcasts to inform all other routers of the internetwork configuration
- Allow routers to perform broadcasts whenever they detect a change in the internetwork configuration

Figure 7-5 shows the contents of the IPX RIP packet.

Operation	2 bytes
Network Number	4 bytes
Hop Count	2 bytes
Tick Count	2 bytes
Up to 50 sets	

Figure 7-5. IPX-RIP Packet

When RIP packets are encapsulated in the data field of the IPX packet, the packet type field is set to one to indicate that the data field contains RIP information. The destination and source socket fields are set to 453 (hexadecimal) to indicate that the source device is sending a request to the routing process on the destination node via the RIP socket number. This procedure is called peer-to-peer processing. The operation field of an IPX packet indicates whether the packet is a request or response. Request packets can be specific (about specific networks) or general (about all networks on an internetwork); response packets can be specific, general, or informational (whenever a router is first started, shut down, or becomes aware of a routing change in the internetwork).

Routers constantly exchange information with each other to ensure that their routing tables reflect up-to-date changes. IPX routers perform RIP broadcasts at the following times:

- Initial broadcasts of directly connected network segments
- Initial requests to receive routing information from other routers
- Periodic broadcasts (every 60 seconds) of the current list of active network numbers
- Broadcast of changes in the internetwork configuration
- Final broadcast when a router or routing device is brought down

Each RIP broadcast is a local broadcast, addressed so that it is not passed on by the routers that receive it. All Novell routing devices use Split Horizon to advertise RIP and SAP information.

Service Advertising Protocol (SAP)

The Service Advertising Protocol (SAP) lets service-providing nodes advertise their services and internetwork addresses. Examples of service-providing nodes are file servers, print servers, queue servers, and so forth. Through the use of SAP advertisements, SAP agents build and maintain a database of internetworking service information known as a server information table. As servers are initialized, they use SAP to advertise their services. When servers are taken down using the down command, they use SAP to indicate that their services are no longer available.

Each client on the network must first obtain the address of a server before it can access any network services. Clients build their own cache tables for up to seven servers. The SAP packet structure allows the following functions:

- A workstation request for the name and address of the nearest server of a certain type
- A general request by a router for names and addresses of all servers or servers of a certain type
- A response to either a nearest server request or a general request
- A 60-second periodic broadcast by servers or routers
- A broadcast of changed server information

As with RIP packets, the SAP packet is encapsulated in the data fields of the IPX packet and the packet type field is then set to 4 to indicate that the data field contains SAP information. Also as with RIP, the destination and source socket fields are set to 452 (hexadecimal) to indicate peer-to-peer processing. Instead of the network number entries found in a RIP packet, SAP response packets can contain from one to seven server entries, whereas SAP request packets contain only the operations field and server type field.

SAP packets contain an operation field that identifies if the packet is a request (general or specific) or a response. Two other operations are Get Nearest Server request or response. The SAP packet also identifies the type of service (such as print queue, file server, job server, print server, archive server, including unknown) and includes the server name with network, node, and socket address and number of hops to the server.

Figure 7-7 shows the contents of the IPX SAP packet.

Operation	2 bytes
Service Type	2 bytes
Server Name	48 bytes
Network Address	4 bytes
Node Address	6 bytes
Socket Address	2 bytes
Hops to Server	2 bytes
Maximum of 7 sets	

Figure 7-6. IPX-SAP Packet

SAP broadcasts are similar to RIP broadcasts and include informing the network that a new server is available, general requests for information, responses to general requests, and 60-second broadcasts containing updated SAP information using split horizon.

Broadcast Techniques

RIP and SAP broadcasts are sent out from the routers and servers using the best route algorithm and split horizon techniques.

Best Route Algorithm

The best route algorithm for Novell routing devices states: If two paths exist, the best route is the one with the least number of ticks. A tick is 1/18th of a second. For local LAN segments with more than 1 MB of bandwidth for each second of bandwidth, the router assumes that the delivery time is one tick. Figure 7-7 illustrates an example of a best route algorithm.

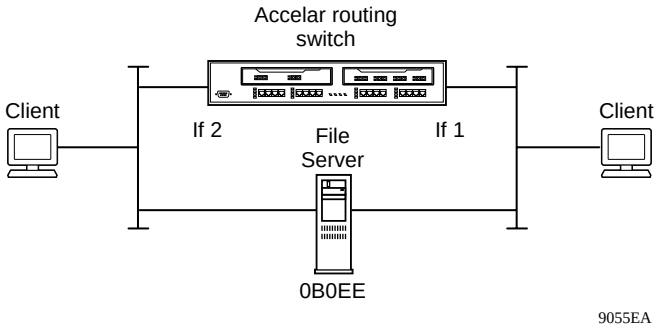

Figure 7-7. Best Route Algorithm Example

If the number of ticks is equal, the best route is the one with fewer hops. A hop is counted each time the packet passes through a router.

If both ticks and hops are equal, the router chooses the first route. When responding to RIP requests and sending informational RIP broadcasts, routers advertise only the best route.

Split Horizon

When routers are initialized, they build a routing table based on their directly connected interfaces. The file server also builds a routing table, and all devices send routing information and service advertising updates using the best route algorithm and split horizon techniques. Split horizon helps reduce bandwidth and speeds up information distribution by advertising only the routing or service information obtained from other interfaces. Information learned from an interface is never advertised out that same interface.

In Figure 7-8, split horizon is enabled between router A and network 4 to eliminate redundant RIP and SAP traffic. If split horizon was disabled on that interface, router B would learn all about network 3 and router C would learn all about network 2.

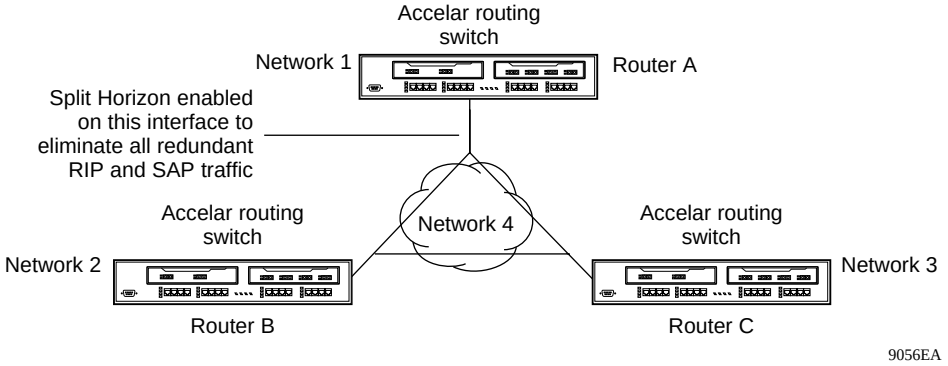

Figure 7-8. Split Horizon Enabled

NetBIOS

The network basic input/output system (NetBIOS) is a connection-oriented communication protocol jointly developed by Sytek Incorporated and IBM. Novell NetWare provides a NetBIOS emulation package that runs on top of the packet exchange protocol (PEP), which provides a service similar to SAP. After processing NetBIOS calls, PEP sends them to IPX. In the NetWare environment, NetBIOS communicates by establishing a logical connection between two NetBIOS-defined names. NetBIOS sets up a logical channel for higher-layer protocols to use for communication.

Figure 7-9 shows a typical environment supporting NetBIOS. When the NetBIOS client wants to make a logical connection to the NetBIOS application running on the server, the client sends a NetBIOS name propagation packet (type 14 hexadecimal).

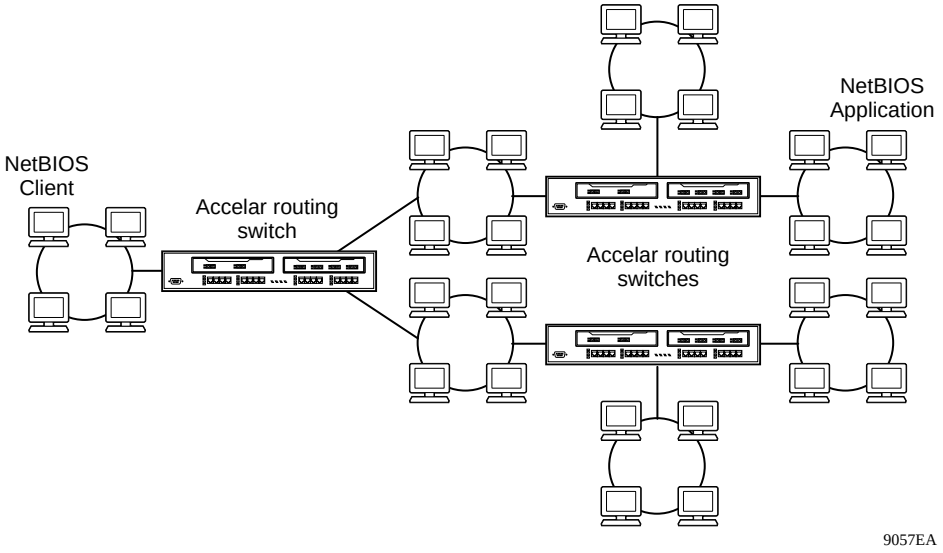

Figure 7-9. NetBIOS Support in a NetWare Environment

Intervening routers perform the following tasks:

- Examine the transport control field to determine the number of routers traversed and the number of network number fields that contain information.
- Compare the network number entries to the network number of the segment where the packet was received. If there is a match, the packet is discarded. If not, the router places the network segment address where the packet arrived into the next available network number field.
- The router increments the transport control field and sends the packet to all directly connected networks not included in the network number fields.
- When the name propagation packet reaches the application server, the server stores the required name and address information and issues a connection number to the client.
- All further communications between the client and server use a standard IPX header with header socket numbers set to 455 hexadecimal to indicate that the packet contains NetBIOS data.

Static Routes

You can configure a static route on each logical IPX interface that the router supports. Using static routes allows you to:

- Direct all IPX traffic destined to a given network to an adjacent host.
- Reduce routing traffic by disabling the RIP supply function on a subset of attached interfaces configured with static routes.
- Control the size of the routing tables.

A static route remains in the routing table until it is deleted or until the router learns a better path based on the selected routing method (hops or RIP ticks).

Static SAP Services

The ability to statically configure services allows services to be advertised locally without the SAP traffic that would be required to learn about the service. A static service is listed in the router SAP table if a route to the network associated with the service exists in the route table. The network associated with the static service must be reachable from the router. A static service remains in the SAP table until you delete it or until the network associated with it becomes unreachable.

IPX Default Static Routes

By default the Accelar routing switch does not contain a default route address in its route table. Without a default static route, the packet is sent to the switch CPU, where it is dropped unless otherwise configured. You can create a default route when adding a static route by using the value of 0xFFFFFFFE as the target network field. IPX uses the reserved network address 0xFFFFFFFE for default routes.

Appendix A
Port Numbering and MAC Address Assignment

This appendix discusses how ports are numbered on the chassis, as well as how MAC addresses are assigned to MAC entities in the Accelar 1000 Series routing switch.

Port Numbering

Some screens contain fields for selecting ports. A port number includes the slot location of the I/O module in the chassis, as well as the port's position in the I/O module. In the Accelar 1200 and 1250 switches, slots are numbered from top to bottom. Figure A-1 shows slot numbering for the Accelar 1200 switch.

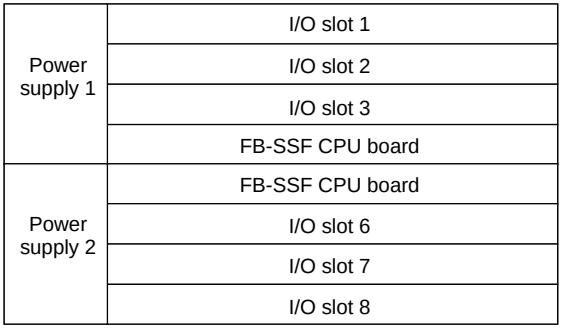

Figure A-1. Accelar 1200 Slots

In the Accelar 1100/1150 switch, the modular slots are slots 1 and 2 and the internal ports belong to slot 3. Figure A-2 shows how slots on an Accelar 1100 chassis are numbered from left to right.

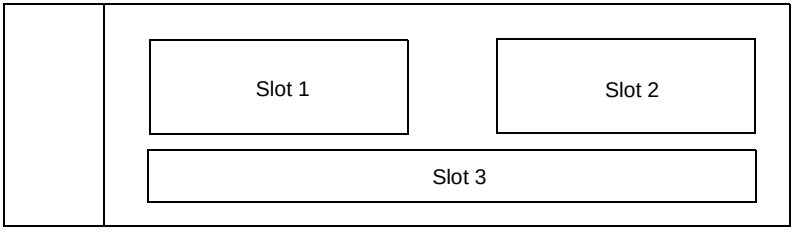

Figure A-2. Accelar 1l00 Slots

Ports in the chassis and in all modules are numbered from left to right. For example, the second port in an I/O module located in slot 1 is 1/2. Figure A-3 shows port numbering in modules.

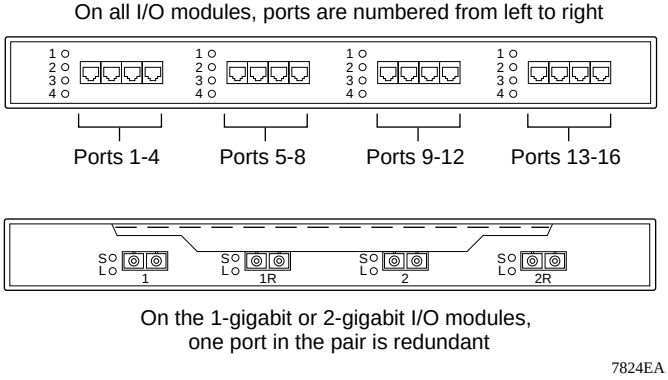

Figure A-3. Port Numbering on I/O Modules

In the Accelar 1050/1051 switch, the 10/100 Mb/s ports are seen as being in slot 3 and the Gigabit port is considered to be in slot 1, with 1R as the redundant port in an Accelar 1051 switch. Figure A-4 illustrates slot and port numbering in an Accelar 1051 switch.

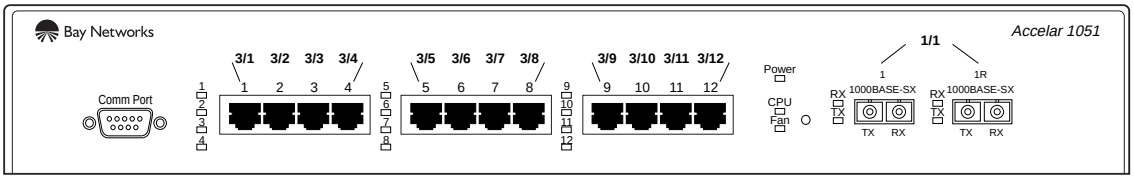

Figure A-4. Slot and Port Numbering on the Accelar 1050/1051 Switch

Use the slot and module examples in the figures as guides when you need help selecting ports in Accelar Device Manager.

MAC Address Assignment

Understanding how MAC addresses are assigned is important when defining static ARP entries for IP addresses in the routing switch and when using a network analyzer to decode network traffic.

Base MAC Address

A flash memory device holds a unique 48-bit base MAC address for the routing switch. For the Accelar 1200 or 1250 chassis, the flash device is in the main chassis. For the Accelar 1100 or 1150 chassis and the Accelar 1050/1051 switch, the flash device is on the main board with the fixed ports.

For a given routing switch, the base MAC address will look like:

```
xx xx xx yy yy 00
```

where:

xx xx xx is the IEEE organization identifier (for example, 00 0E 16).

yy yy is unique to the routing switch.

On switches with debug Ethernet ports, the base MAC address is used by this port.

The general form for a MAC addresses used by a particular routing switch is:

```
xx xx xx yy yy zz
```

where:

xx xx xx is the IEEE organization identifier (for example, 00 0E 16).

yy yy is unique to the routing switch.

zz is user specific.

From the general form, it is easy to see that each Accelar 1000 Series routing switch is assigned a block of 256 48-bit MAC addresses from xx xx xx yy yy 00 through xx xx xx yy yy FF.

Physical MAC Addresses

Physical MAC addresses are addresses assigned to the physical interfaces or ports visible on the device. The physical MAC addresses are used in the following types of frames:

- Spanning Tree Protocol BPDUs sent by the routing switch
- Frames to or from an isolated routing port's physical interface

BPDUs are sent using the physical MAC address as the source because identifying which physical port sent the BPDU is critical to how the Spanning Tree Protocol works. For isolated routing ports, the IP address is associated with the physical interface, so the physical MAC address is associated with the IP address.

The last byte of the MAC address (zz in the general form) for a physical interface depends on the slot and port number for the given interface. The basic scheme is that each slot is allocated 16 physical MAC addresses. If a board has fewer than 16 ports, some MAC addresses are unused. Table A-1 lists the value for the last byte of the MAC address based on the slot and port number.

Table A-1. Last Byte of Physical MAC Address

Slot	Port															
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**15**
1	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
2	10	11	12	13	14	15	16	17	18	19	1A	1B	1C	1D	1E	1F
3	20	21	22	23	24	25	26	27	28	29	2A	2B	2C	2D	2E	2F
4	Not applicable—contains SSF module															
5	Not applicable—contains SSF module															
6	30	31	32	33	34	35	36	37	38	39	3A	3B	3C	3D	3E	3F
7	40	41	42	43	44	45	46	47	48	49	4A	4B	4C	4D	4E	4F
8	50	51	52	53	54	55	56	57	58	59	5A	5B	5C	5D	5E	5F

Slots 4 and 5 do not have any MAC addresses assigned to them. Depending on which switch you are operating, slots 4 and 5 may not be present or hold the SSF modules.

For example, a switch with the base MAC address 00 0E 16 11 00 00 has a physical MAC address for slot 3 port 6 (port 3/6) of 00 0E 16 11 00 25. This MAC address is seen as the source MAC address for any BPDUs sent out of this port. If port 3/6 is configured as an isolated routing port, ARP requests sent to the IP address of the isolated routing port will return this MAC address.

Virtual MAC Addresses

Virtual MAC addresses are the addresses assigned to VLANs. A virtual MAC address is assigned to a VLAN when it is created. The MAC address for a VLAN IP address is the virtual MAC address assigned to the VLAN.

The range for the last byte of the virtual MAC addresses in hex is 81 through FF, that is, the most significant bit of the last byte, is set to 1.

A virtual MAC address is assigned when a VLAN is created. Since the Default VLAN (VLAN ID 1) is always created, the last byte of the MAC address for VLAN 1 is always 81. For other VLANs, the MAC address assigned can be found in Device Manager (VLAN > VLAN > Advanced) or through the Run-Time CLI (*show vlan info advance* command).

Index

Numbers

802.1Q tagging, 3-9

A

Accelar Configuration page, 2-10
Accelar Device Manager, 2-8, 3-22
Accelar routing switch overview, 1-1
Accelar software, 1-2
Accelar switch models, 1-1
Accelar VLAN Manager, 2-8, 3-22
accept policies, 4-21
access policies, 2-13
address/mask pair, 4-3
addresses, multicast, 5-2
announce policies, 4-21
area border router, OSPF implementation, 4-19
area, OSPF, definition of, 4-18
ARP
  function of, 4-6
  table window, 4-7
AS boundary router, OSPF implementation, 4-19
ASICs, identifying, 1-3

B

backup designated router, 4-20
base MAC address, A-3
Best Effort priority mode, 3-13
best route algorithm, 7-10
boot configuration, bypassing, 2-4
boot flash, 2-4
Boot Monitor CLI, 2-9
boot monitor image, 2-7
boot sequence, 2-1
BootP, 4-8
BootP/DHCP
  relay, 4-8
  servers, 4-10
Bootstrap Protocol. *See* BootP
BPDUs, 3-21
Bridge Protocol Data Units. *See* BPDUs
bridges, spanning tree, 3-10
broadcast domains, 3-1
brouter ports, 3-16, 4-3, 4-5

C

CIDR (classless interdomain routing) address, 4-3
CLI
  access, 2-12
  Boot Monitor, 2-9
  for network management, 3-22
  reference for, 1-1
  Run-Time, 2-9
collision domains, 3-1
command line interface. *See* CLI
compatibility, hardware and software, 1-2
configuration
  default, 2-3, 3-17
  files, 2-7
  loading, 2-3
connectionless datagram protocol, 7-1

D

datagram, 7-1

default route, IPX support, 7-7

default static routes, 4-6

default VLAN, described, 3-16

designated router, 4-20

destination filters, 6-2

Device Manager, 1-1, 2-8, 3-22

DHCP relay operation, 4-8

distance vector protocols
- DVMRP, 5-6
- RIP, 4-15

DVMRP
- Accelar support of, 5-10
- Distance Vector Multicast Routing Protocol
- multicast tree, 5-7
- reverse path forwarding, 5-7
- routes, 5-9
- routing table, 5-9
- source networks, 5-6, 5-8

Dynamic Host Configuration Protocol. *See* DHCP

dynamic routing, 7-7

E

education services, xx

executables, 2-6

explicit tagging, 3-9

F

factory default load, 2-3

filters
- destination, 6-2
- global, 6-3
- IP, 6-1
- source, 6-2

frames, tagging, 3-9

G

global filters, 6-3

H

hop count, 4-15

host groups, multicast, 5-1

I

IGMP
- Accelar support of, 5-5
- characteristics, 5-3
- host leave message, 5-4
- host reports, 5-4
- Internet Group Management Protocol. *See* IGMP
- proxy, 5-6
- queries, 5-3
- snooping, 5-5

IGP, 4-16

implicit tagging, 3-10

interior gateway protocol. *See* IGP

internal router, 4-19

Internet Data Packet (IDP) format, 7-2

internetwork, 7-2

intranode addressing, 7-2

IP accept policies, 4-20

IP address
- definition of, 4-1
- in dotted decimal notation, 4-1

IP announce policies, 4-21

IP filters
- Accelar support of, 6-1
- characteristics, 6-1
- configurable actions, 6-3
- default action, 6-1
- matching criteria, 6-3
- port actions, 6-4

IP routing, 4-1

IP routing protocols supported, 4-1

IPX
- Accelar support of, 7-1
- basic network, 7-2
- default static routes, 7-14
- file servers, 7-3
- frame formats, 7-1
- header, 7-5
- internetwork addressing, 7-2

NetBIOS, 7-12
network-level services, 7-7
packets, 7-4
RIP, 7-7
RIP packets, 7-8
SAP, 7-9
SAP packets, 7-10
socket numbers, 7-2
static routes, 7-14
static service, 7-14
isolated routing ports, 4-3

L

L2 Read-Write access, 2-11
L3 Read-Write access, 2-11
local segments, 3-1
log files, 2-7

M

MAC address assignment, A-3
Magic Cookie, 4-8
management access, 2-13
management information base. *See* MIB
members, VLAN designations, 3-3
metric, defined, 4-15
MIB, 3-22
MLT
defined, 3-17
examples, 3-18
requirements, 3-17
multicast stream, 5-4
multicasting
addresses for, 5-2
comparing routes, 5-10
leaf network, 5-10
shortest-path tree, 5-10
Multi-Link Trunking. *See* MLT

N

NetBIOS
defined, 7-12
name propagation packet, 7-12
name service, 4-11
network basic input/output system. *See* NetBIOS
Network Information Center (NIC), 4-1
non-tagged ports, 3-9
not so stubby areas. *See* NSSAs
Novell networks, 7-3
Novell, Inc., 7-1
NSSAs, 4-19

O

OSPF
area, definition of, 4-18
database synchronization, 4-17
features
link state protocol, 4-17
shortest-path tree, 4-17
implementation
area border router, 4-19
AS boundary router, 4-19
backup designated router, 4-20
designated router, 4-20
other router, 4-20
OSPF. *See* Open Shortest Path First, 4-16
other router, definition, 4-20

P

passwords, Accelar, 2-11
PCMCIA
card description, 2-6
file system, 2-4
permanent host group, 5-1
physical MAC address, A-4
physical router interfaces, 4-3
PIDs
defined, 3-6
reserved, 3-6
policies
accept, 4-21
announce, 4-21
definition of, 4-21
port locking, 2-14
port mirroring, 3-23

ports
isolated routing, 4-3
locking, 2-14
numbering, A-1
Spanning Tree FastStart, 3-12
prioritization, traffic, 3-13
priority mode, 3-13
product support, xx
Protocol Identifiers. *See* PIDs
protocol-based VLANs supported, 3-5
proxy ARP, 4-7
publications
Bay Networks related, xvii
external related, xviii

R

RARP, 3-6, 4-12
RARP VLANs, 4-12
Read-Only access, 2-11
Read-Write access, 2-11
Read-Write-All access, 2-11
Reverse Address Resolution Protocol. *See* RARP
RFCs, xviii
RIP
described, 4-15
in IPX networks, 7-7
RMON, 3-22
router ARP table, 4-7
router interfaces, 4-3
Routing Information Protocol. *See* RIP
routing tables, flushing, 4-8
Run-Time CLI, 2-9
run-time image, 2-2, 2-6

S

SAP
advertisements, 7-9
definition, 7-9
script files, 2-7
security, 2-11
Service Advertising Protocol. *See* SAP
severity codes, 3-24
severity levels, syslog, 3-25
Silicon Switch Fabric. *See* SSF
SNMP community strings, 2-12
SNMP traps, 3-23
socket numbers, 7-2
source filters, 6-2
spanning tree bridges, 3-10
spanning tree groups, 3-10, 3-11
Spanning Tree Protocol. *See* STP
Split Horizon technique, 7-8, 7-11
SSF module, boot image, 2-2
static routes
creating a default, 4-6
support, 7-7
STP, 3-10, 3-11
stub areas, 4-19
subnet mask, function of, 4-2
subnets, definition of, 4-2
supernet address, 4-3
supernets, definition of, 4-2
support, Bay Networks, xx
syslog
definition, 3-24
severity levels, 3-25
syslogd daemon, 3-24
system flash, 2-5

T

tagged ports, 3-9
technical support, xx
trace logs, 2-7
traffic filters
actions, 6-4
defined, 6-2
traffic prioritization, 3-13
traffic queues, 3-13
transient host group, 5-1

transit network, 3-15

U

UDP, 4-11
UDP broadcast forwarding, 4-11
unassigned VLAN, described, 3-16
UNIX messages, 3-24
User Datagram Protocol. *See* UDP

V

virtual IP addresses, 4-12
virtual local area networks. *See* VLANs
virtual MAC address, A-5
virtual router ID, 4-13
virtual router interfaces, 4-3, 4-4
Virtual Router Redundancy Protocol. *See* VRRP
VLAN Manager, 2-8, 3-22
VLANs
  default, 3-16
  definition, 3-2
  performance rules, 3-14
  policy-based, 3-3
  port-based, 3-2
  protocol-based, 3-5
  rules, 3-14
  source IP subnet-based, 3-7
  source MAC-based, 3-4
  tagging, 3-8
  unassigned, 3-16
VRRP, 4-12
VRRP routers, backup and master, 4-14

W

Web page, 2-10
Web window, 2-13

X

Xerox Network System (XNS), 7-1